Stress-driven strategic games in cancer

Selahattin Can Ozcan[1,*]

[1]Department of Physiology and Cellular Biophysics, Columbia University, New York City, NY, USA
[*]Corresponding author: so2716@cumc.columbia.edu

ABSTRACT

Tumor cells face chronic genotoxic, metabolic, hypoxic, and immune stress that shapes their evolution. While stress-response pathways are well characterized molecularly, cancer biology lacks a predictive framework for how cells select among alternative adaptive strategies and how these selections interact to produce tumor-level behavior. We propose that evolutionary game theory, previously applied to cooperation in cancer, should be extended to position stress adaptation itself as the organizing principle of tumor evolution. In this framework, stress-adaptive strategies constitute frequency-dependent games whose payoffs depend on population composition. We introduce a three-level distinction between cell states (transcriptional snapshots), game states (local configurations of stress and neighbor composition that define the active payoff structure), and cell strategies (conditional behavioral policies mapping game states to fitness-relevant outputs). This perspective explains the maintenance of intratumor heterogeneity through frequency-dependent selection, the reversibility of resistance through bet-hedging dynamics, and therapy resistance as an equilibrium outcome rather than genetic inevitability. Integrating insights from single-cell genomics, spatial profiling, and lineage tracing, we outline testable predictions and experimental approaches for measuring payoff structures. Therapeutically, the framework suggests exploiting adaptive trade-offs, restricting phenotypic plasticity, and reshaping competitive landscapes. Re-framing cancer as an evolving game of stress adaptation provides a unifying structure for predictive oncology.

Introduction

Tumor cells are exposed to a wide range of stresses throughout cancer progression, including DNA damage arising from replication stress, oxidative stress generated by altered metabolism, nutrient limitation, hypoxia, proteotoxic stress from high biosynthetic demand, immune-mediated cytotoxicity, and mechanical constraints imposed by tissue architecture. Importantly, these stresses are rarely acute or uniform; instead, they are chronic, spatially heterogeneous, and often sublethal. Such conditions are sufficient to exert strong selective pressure without causing population collapse, favoring cellular strategies that enable persistence under sustained stress rather than maximal short-term proliferation.

Across cancer types, stress does not simply accumulate as damage but actively shapes cellular behavior through engagement of conserved stress response programs[1]. Persistent genotoxic stress activates the DNA damage response (DDR), often selecting for tolerance or checkpoint adaptation rather than faithful repair. Oxidative stress generated by altered metabolism engages antioxidant and redox-buffering programs, frequently mediated by the oxidative stress response (OSR), exemplified by NRF2-mediated transcriptional programs that upregulate glutathione synthesis, NADPH production, and ROS-detoxifying enzymes. Proteotoxic stress arising from high translational and biosynthetic demand activates the unfolded protein response (UPR) and the heat shock response (HSR), stabilizing protein homeostasis at the cost of altered signaling and metabolic states. Hypoxic stress, driven by poor and fluctuating perfusion, engages the hypoxic stress response (HySR), rewiring metabolism, angiogenesis, and survival pathways through HIF-dependent and HIF-independent mechanisms. These adaptations recur across genetically distinct tumors, indicating convergent evolution under shared stress constraints rather than idiosyncratic responses to damage (**Fig. 1A**).

Crucially, the stresses shaping cancer evolution arise from both cell-autonomous and non–cell-autonomous sources, which together define the selective landscape experienced by tumor cells. Cell-autonomous stresses originate from intrinsic properties of malignant cells, including oncogene activation, chromosomal instability, dysregulated metabolism, and aberrant biosynthesis, and continuously generate replication stress, proteotoxic burden, and redox imbalance. For example, MYC-driven tumors experience constitutive replication stress and nucleotide depletion[2], while KRAS-mutant cells face chronic oxidative stress from altered mitochondrial metabolism[3]. In contrast, non–cell-autonomous stress refers to stressors imposed by the tumor microenvironment, arising from interactions with neighboring cells, extracellular matrix, tissue architecture, and systemic physiology. These stresses are externally generated, context dependent, and frequently shared across local cell populations, making them powerful selective forces that operate at the level of niches rather than individual cells (**Fig. 1B**).

Non–cell-autonomous stresses encompass multiple, interrelated classes that collectively constrain cellular behavior. Mechanical and physical stresses arise from solid stress, compression, shear forces, and extracellular matrix stiffness imposed by

tumor growth within confined tissue spaces, altering cytoskeletal organization, nuclear integrity, and proliferative capacity. In pancreatic ductal adenocarcinoma, for instance, desmoplastic stroma generates solid stress exceeding 75 mmHg[4], compressing blood vessels and creating nutrient gradients that stratify metabolic phenotypes across the tumor mass. Metabolic and nutrient stresses reflect competition for shared resources, including glucose, amino acids, lipids, and oxygen, driven by abnormal vasculature, poor perfusion, and consumption by neighboring tumor and stromal cells. Hypoxic stress emerges from vascular insufficiency and fluctuating blood flow, producing chronic or cycling hypoxia that reinforces metabolic, proliferative, and survival trade-offs. Immune-mediated stress is imposed through cytotoxic killing, inflammatory cytokine signaling, and immunoediting, often acting as a persistent and sublethal pressure rather than an all-or-none elimination force. Studies in melanoma and NSCLC have demonstrated that chronic interferon-gamma exposure selects for JAK1/JAK2 loss-of-function mutations[5], representing adaptation to sustained immune pressure. Finally, stromal and ecological stresses arise from paracrine signaling, matrix remodeling, spatial crowding, and clonal competition, constraining cellular behavior through collective and spatial effects that cannot be reduced to cell-intrinsic pathways. Additionally, neural innervation recruited through neoneurogenesis[6] and the intratumoral microbiota, which can directly metabolize chemotherapeutics and reshape immune polarization[7,8], introduce further ecological dimensions to the stress landscape whose game-theoretic implications are discussed below.

Taken together, the diversity of stresses confronting tumor cells should not be viewed as independent insults requiring independent responses. Rather, they form an interconnected stress landscape in which each stress constrains the solutions available for others. A cell investing heavily in oxidative stress defense may lack the resources for rapid proliferation, while a cell that suppresses immune recognition may sacrifice metabolic flexibility. These constraints mean that stress does not merely damage cells but forces them to make choices, to allocate finite resources among competing adaptive programs, to specialize or remain flexible, to cooperate with neighbors or compete against them. The observation that genetically diverse tumors converge on similar phenotypic architectures across patients suggests that these choices are not random but reflect structured solutions to recurring stress problems. Understanding how cells navigate these decisions, and how the choices of individual cells interact to produce population-level outcomes, requires a formal framework capable of capturing the interdependence of adaptive strategies within heterogeneous, evolving populations.

This Review adopts a normative stance, namely that stress adaptation, rather than signaling pathways or mutations, should serve as the primary organizing principle for understanding tumor evolution and designing therapeutic interventions. We distinguish three modes of academic speech, descriptive (what is observed), interpretive (what it means), and normative (what should change), and intentionally operate in all three. What follows is not a neutral survey of the literature but a framework-defining argument, grounded in evidence and designed to generate testable predictions.

Throughout this Review, we use the terms of evolutionary game theory with specific operational definitions that distinguish three levels of description. A *cell state* is a snapshot of a cell's gene expression at a given moment, reflecting what the cell is currently doing in response to its environment (environment → transcriptional response → snapshot); cell states are what single-cell RNA sequencing resolves directly. A *game state* is the local configuration of stresses, resources, and neighbor strategy compositions that collectively define which payoff structure is currently in effect for a cell (microenvironment + neighbor composition → payoff structure); two cells in identical transcriptional states can face different game states if their neighborhoods differ, and therefore adopt different strategies. A *cell strategy* is a conditional behavioral policy that determines how a cell allocates resources and signals to neighbors given the game state it finds itself in, producing fitness consequences that depend on what other cells around it are doing (game state + regulatory architecture → behavioral output → neighbor-dependent fitness consequences). Strategies are thus not synonymous with cell states: a cell state tells us where a cell is, the game state tells us what move it should make, and the strategy is what it does—the mapping from game state to behavioral output. Recurrent transcriptional programs (e.g., proliferative, glycolytic, drug-tolerant persister, immune-evasive) provide the observable signatures from which strategies can be inferred, but the same transcriptional state may correspond to different strategies under different game states. A *payoff* is the per-capita net growth rate of cells executing a given strategy, measured over a defined interval. *Frequency dependence* means that a strategy's payoff changes as the relative abundance of strategies in the local population shifts. The *game* is the mapping from population strategy composition to the vector of strategy-specific payoffs—i.e., the payoff matrix—which may itself change when the stress environment changes (a game switch).

From Stress Response to Stress Strategies

From pathways to strategies

Cancer biology has traditionally framed stress adaptation in terms of discrete signaling pathways activated by specific insults. DNA damage activates checkpoint signaling, hypoxia induces transcriptional programs, and nutrient limitation engages metabolic sensors. While this framework has yielded deep mechanistic insight, it implicitly assumes that stress responses are linear and deterministic. Ecological perspectives, particularly Lotka-Volterra competition models[9] and predator-prey dynamics applied to tumor-immune interactions[10], have enriched this picture by demonstrating that population-level outcomes depend on density-dependent interactions and resource competition, not merely cell-intrinsic programs. Game theory builds on these

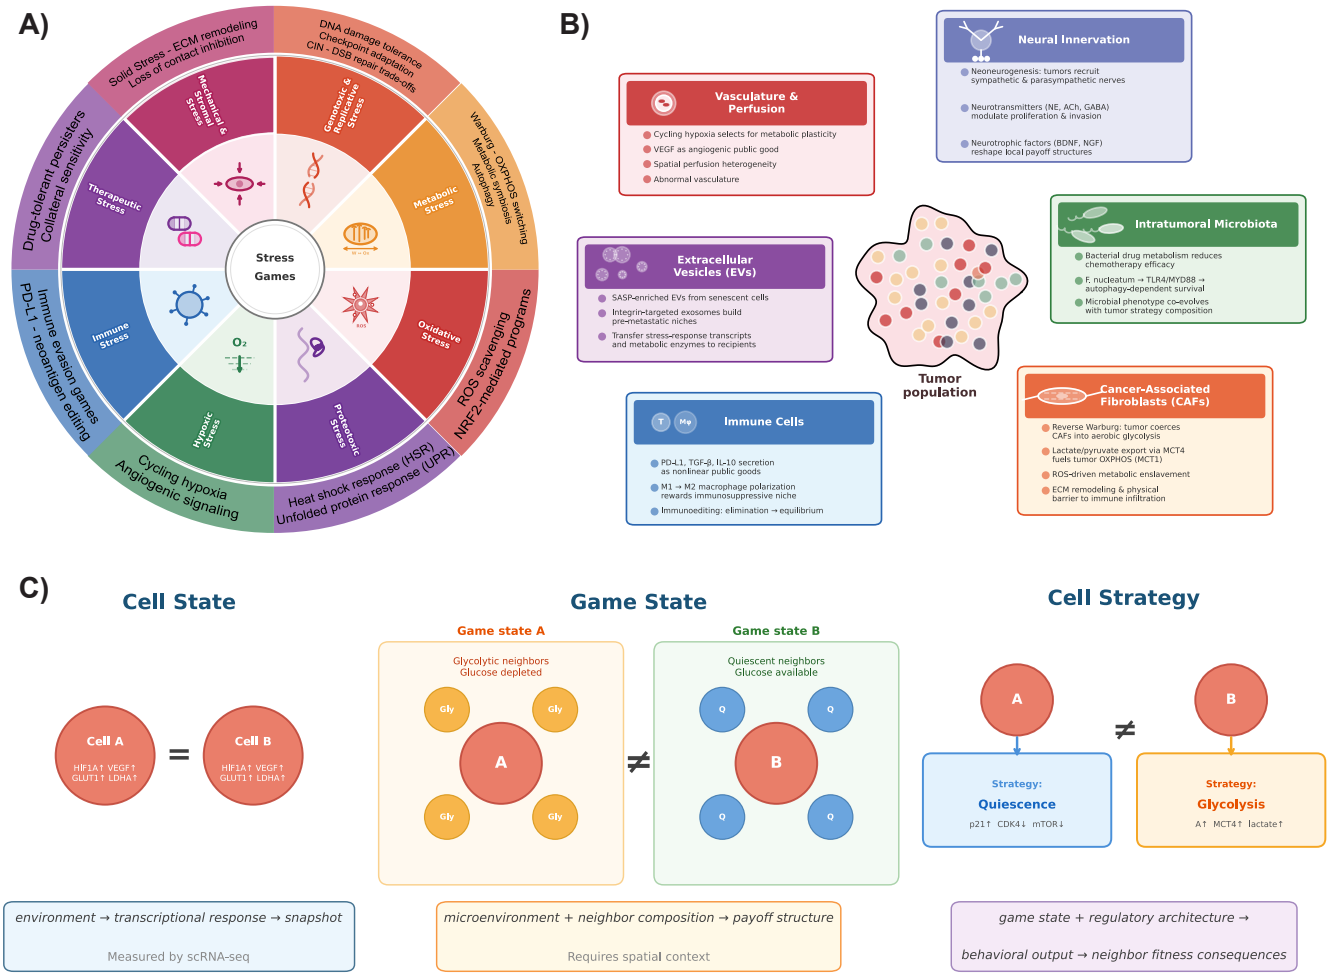

Figure 1. The cancer stress landscape and its microenvironmental players. (A) Eight cancer-associated stress domains and their game-theoretic strategies. The circular wheel organizes stresses into genotoxic & replicative, metabolic, oxidative, proteotoxic (UPR/HSR), hypoxic, immune, therapeutic, and mechanical & stromal categories, each paired with the dominant adaptive strategies they select for. Central to this framework is the concept that these stress domains do not operate independently but jointly define the selective landscape in which tumor cell strategies compete and coexist. **(B)** The tumor microenvironmental game and the extracellular players that reshape stress-adaptive payoffs. Six non–cell-autonomous players, vasculature & perfusion, extracellular vesicles (EVs), immune cells, neural innervation, intratumoral microbiota, and cancer-associated fibroblasts (CAFs), surround the central tumor population and modulate the payoff matrix through bidirectional signaling. Each player box summarizes key mechanisms by which that compartment alters the relative fitness of competing tumor strategies, highlighting how microenvironmental context shapes evolutionary game outcomes beyond cell-intrinsic programs. **(C)** The distinction between cell state, game state, and cell strategy. Two genetically identical cells (Cell A and Cell B) occupy the same transcriptional state, yet adopt different strategies because they face different game states. Cell A, surrounded by glycolytic neighbors in a glucose-depleted niche, adopts a quiescence strategy, whereas Cell B, surrounded by quiescent neighbors with glucose available, adopts a glycolytic strategy. Cell states are transcriptional snapshots measured by scRNA-seq; game states are local configurations of stress and neighbor composition that define the active payoff structure; and cell strategies are conditional behavioral policies that map game states to fitness-relevant outputs. This hierarchy illustrates that strategy cannot be inferred from transcriptional state alone but requires knowledge of the local game state.

ecological foundations by adding a critical dimension, namely that the fitness consequences of any cell's adaptive strategy depend not only on resource availability and population density but on the specific strategies adopted by other cells in the population[11,12] (**Box 1**).

The application of evolutionary game theory to cancer was pioneered by Axelrod and colleagues[11], who recognized that tumor cells engage in both competitive and cooperative interactions that determine fitness in frequency-dependent ways. This foundational insight was synthesized comprehensively by Archetti and Pienta[13], who demonstrated that cooperation among cancer cells, mediated by the secretion of diffusible growth factors, immune-suppressive cytokines, and microenvironment-modifying metabolites, constitutes a multiplayer public goods game rather than a simple pairwise interaction. Crucially, they showed that the relationship between the number of cooperating cells and the collective benefit is typically nonlinear, following a sigmoid function rather than a linear sum; this nonlinearity fundamentally alters whether cooperation can be maintained at intermediate frequencies, enabling stable polymorphic equilibria between cooperators and defectors that linear models would not predict. Empirical validation came from Archetti and colleagues' demonstration[12] that IGF-II production in pancreatic neuroendocrine tumors exhibits classic nonlinear public goods game dynamics, with producer and non-producer cells coexisting through frequency-dependent selection at proportions consistent with theoretical predictions. However, cooperation represents only one class of interactions within tumors. The game-theoretic literature on cancer has focused predominantly on the cooperator–defector axis, asking why cells produce shared resources when non-producers can free-ride. Here, we extend this framework beyond cooperation to encompass the full spectrum of stress-adaptive strategies, including metabolic rewiring, DNA damage tolerance, immune evasion, and phenotypic plasticity. Where game-theoretic models of cooperation ask why cells secrete shared resources despite the threat of cheaters, we pose a broader question: how do cells choose among alternative adaptive solutions to persistent stress, and how do these choices interact to shape tumor evolution? We argue that this framework must be extended beyond cooperation, such that stress itself, rather than cooperation per se, serves as the organizing principle that defines the strategy space and payoff structure of the evolutionary game. Our framework differs from prior game-theoretic models of cancer in two specific respects. First, it identifies stress as the game constructor, the force that defines the strategy space and generates the payoff matrix, rather than treating stress as one of several payoff modifiers within a pre-existing game. Second, it treats the full repertoire of stress-adaptive responses (metabolic, genotoxic, immune, plastic) as a unified strategy set within a single game, rather than modeling each type of interaction as a separate, disconnected game. The result is a framework in which cooperation, competition, and plasticity arise as interconnected consequences of stress-imposed trade-offs rather than as independent phenomena requiring independent models.

Several complementary formalisms have been applied to cancer evolution, each with characteristic strengths and blind spots. Fitness landscape theory[14,15] provides geometric intuition for genotype-fitness mapping but assumes fitness is context-independent, an assumption violated in cancer, where the fitness of any strategy depends on population composition. Adaptive dynamics[16,17] explicitly models frequency-dependent selection but assumes monomorphic resident populations and small continuous mutations, sitting uncomfortably with the large, discrete phenotypic transitions observed in cancer (EMT, persister emergence) that arise through epigenetic mechanisms decoupled from genomic change. Population genetics approaches[18,19] excel at predicting temporal dynamics of resistance mutations but treat resistance as genetically determined and ecologically independent, necessarily concluding that resistance is inevitable given sufficient mutation rates, a prediction that obscures the reversible, frequency-dependent nature of phenotypic resistance strategies operating on far shorter timescales[20] (**Box 2**).

Evolutionary game theory occupies a complementary niche because it was designed to model interactions where fitness is explicitly frequency-dependent. By treating cancer phenotypes as strategic types competing for shared resources, EGT naturally captures the reality that fitness depends on population composition: a proliferative strategy thrives when competitors are scarce but becomes vulnerable when abundant, while a metabolically flexible phenotype gains advantage as others exhaust local resources. The framework accommodates mixed strategies and stochastic phenotypic switching[21] without requiring genetic specification of each strategy, a critical flexibility given incomplete understanding of switching mechanisms. EGT also provides a natural language for therapeutic intervention, because treatments alter payoff matrices (chemotherapy increases proliferation costs)[22], restrict strategy sets (CDK4/6 inhibition blocks cell-cycle progression)[23], or exploit Nash equilibria to drive populations toward favorable configurations. These advantages come with costs. EGT assumes discrete strategy types whereas cancer phenotypes may occupy continuous spaces, and payoff matrices must be inferred from data, a measurement problem harder than determining mutation rates[24]. However, single-cell transcriptomics has lent empirical support to the discrete-strategy approximation by revealing that tumor cells consistently organize into identifiable cell states, recurrent transcriptional programs for proliferation, invasion, drug tolerance, or immune evasion, that persist across patients and cancer types[25,26]. These cell states provide the observable substrate from which strategies can be inferred, although the mapping from cell state to strategy is not one-to-one: the same transcriptional state may implement different strategies depending on the local game state defined by neighbor composition and microenvironmental context.

However, the same stress frequently produces divergent outcomes across tumors or even within the same tumor. DNA damage may induce apoptosis, senescence, quiescence, or continued proliferation through tolerance mechanisms. In BRCA-

deficient cancers treated with PARP inhibitors[27], some cells die through synthetic lethality, others enter senescence, and a subset activates replication fork stabilization pathways[28] to continue proliferating despite persistent DNA damage. Similarly, metabolic stress may favor efficiency through OXPHOS upregulation, scavenging through macropinocytosis[29] and autophagy, or dormancy through mTOR suppression[30]. In melanoma cells deprived of glutamine, transcriptomic and metabolomic profiling reveals at least three distinct adaptive states, including asparagine synthetase upregulation enabling glutamine independence, increased reliance on branched-chain amino acid catabolism, and entry into a quiescent state with minimal biosynthetic activity[31,32]. These observations challenge the notion of a single optimal response and instead suggest the presence of alternative adaptive solutions[33].

Evolutionary game theory provides a language for describing such situations. In this framework, cells do not execute fixed programs but adopt strategies, conditional behavioral policies whose success depends on the game state defined by environmental conditions and the distribution of other strategies in the local population (**Fig. 2A**). Transcriptional cell states provide the observable readout of these strategies but are not identical to them: the same hypoxic cell state may correspond to a glycolytic strategy or a quiescent strategy depending on whether the cell's neighbors are glycolytic (depleting glucose) or quiescent (leaving glucose available). Stress responses become implementations of broader strategic choices rather than endpoints in themselves. A "high glycolysis" cell state, for instance, is not simply a metabolic configuration but reflects a strategic prioritization of ATP generation speed over efficiency, with predictable consequences for lactate production, pH modulation, and competitive interactions with neighboring cells, consequences that depend on the game state, not the cell state alone.

Trade-offs and frequency dependence

A central insight of evolutionary game theory is that fitness is often frequency dependent, meaning the success of a strategy depends on how common it is. This principle has been repeatedly validated in biological systems, from microbial cooperation to host–pathogen interactions.

In cancer, frequency dependence offers a natural explanation for stable heterogeneity. Strategies that prioritize rapid growth may outperform stress-tolerant strategies when resources are abundant, yet become liabilities under stress. Conversely, stress-buffering strategies may carry proliferative costs that limit their dominance. The archetypal example comes from glioblastoma, where rapidly cycling cells with high metabolic demand coexist with slow-cycling, therapy-resistant cells displaying elevated oxidative stress defenses[34]. This interaction exhibits the structure of a Snowdrift (Hawk–Dove) game in which glycolytic cells acidify the microenvironment, creating conditions that glycolysis-adapted cells tolerate better but that disadvantage cells dependent on OXPHOS, while OXPHOS cells consume less glucose, preventing complete resource depletion. Unlike a Prisoner's Dilemma, where defection invariably dominates, the Snowdrift structure ensures that neither strategy can exclude the other, each gains a frequency-dependent advantage when rare (**Fig. 2B**). The coexistence of such strategies is not accidental but emerges from trade-offs inherent to adaptation. Moreover, the benefits of metabolic cooperation in such systems are typically nonlinear, such that the marginal advantage of an additional glycolytic cell diminishes as lactate accumulates and pH drops, while the marginal advantage of an OXPHOS cell increases as glucose competition relaxes[13]. These nonlinear payoff functions generate stable mixed equilibria that would not arise under the linear benefit assumptions of simpler models.

Recent computational work has extended these two- and three-strategy models to higher-dimensional games incorporating proliferative, invasive, resistant, and cooperative phenotypes simultaneously. Spatial evolutionary game models applied to multilayer tumor architectures[22] have revealed that increasing strategic complexity generates novel coexistence regimes and noise-stabilized states not accessible in simpler models, while optimal control approaches framed as Stackelberg games between clinicians and tumor populations[35] demonstrate how environmental perturbations reshape the payoff landscape in ways that can be exploited therapeutically.

Direct evidence for frequency-dependent selection comes from mixing experiments in multiple cancer models. When EGFR-inhibitor-sensitive and EGFR-inhibitor-tolerant lung cancer cells are co-cultured at varying ratios and exposed to drug, the tolerant population expands during treatment but contracts upon drug withdrawal, demonstrating conditional fitness that depends on both the treatment state and the relative abundance of each subpopulation. Crucially, tolerant cells proliferate more slowly than sensitive cells in the absence of drug, explaining their minority status in treatment-naive tumors.

This view reframes tumor heterogeneity as an expected outcome of evolution under constraint rather than as biological noise or experimental artifact. The observation that tumors reproducibly generate similar phenotypic architectures across patients, despite genetic heterogeneity, suggests that these configurations represent evolutionarily stable states shaped by recurring stress landscapes.

Three important boundaries of this framework should be stated explicitly. First, the framework requires frequency dependence, meaning that if a stress response is purely cell-autonomous with no dependence on population composition, for example, a cell's intrinsic UV-damage repair operating independently of its neighbors, then it constitutes a decision problem, not a game. Second, the framework addresses heterogeneity maintained by balancing selection, not heterogeneity arising

from neutral drift in the absence of selection; neutrally drifting clonal diversity, while biologically important, lies outside the explanatory scope of game theory. Third, the discrete-strategy approximation used throughout is empirically motivated by recurrent cell states in single-cell data, which serve as observable proxies for the underlying strategies; however, this approximation imposes limits on resolution. Phenotypic gradients that do not cluster into identifiable states are not well captured by payoff matrix formalism, and, critically, the inference from observed cell states to strategies requires knowledge of the local game state (neighbor composition, microenvironmental context), information that is only beginning to become available through spatial transcriptomics and multiplexed imaging. Continuous-strategy extensions, while mathematically available, remain poorly constrained by current data.

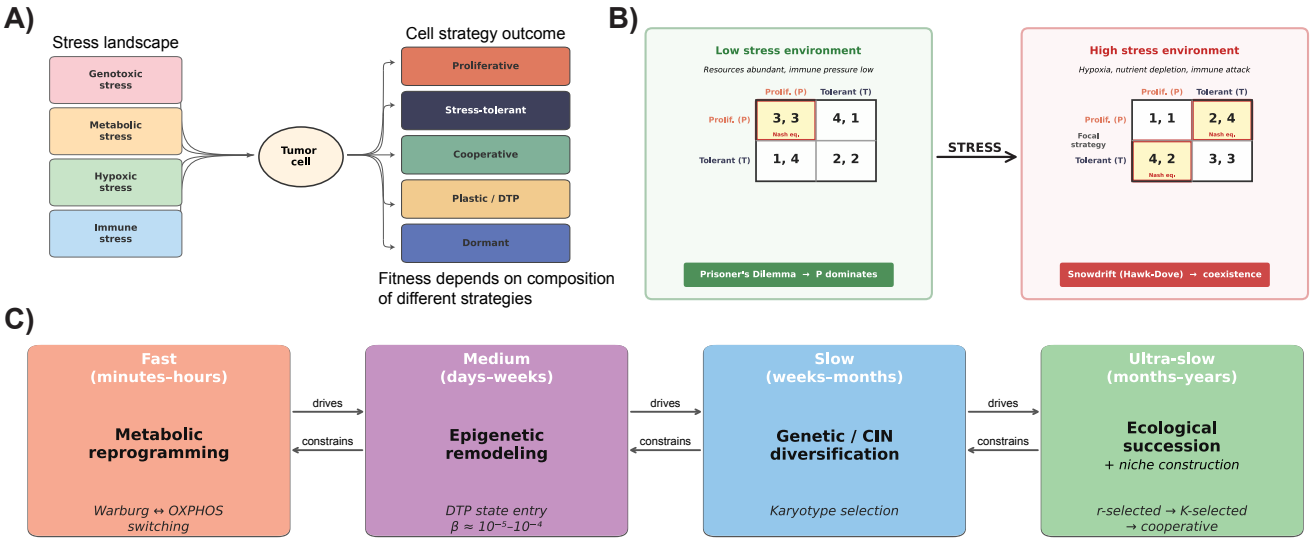

Figure 2. Evolutionary game theory framework for stress-adaptive strategy dynamics. (A) Stress landscape maps to strategy space. Multiple concurrent stresses (genotoxic, metabolic, hypoxic, immune) converge on tumor cells, which respond by adopting one of several discrete strategies, proliferative, stress-tolerant, cooperative, plastic/drug-tolerant persister (DTP), or dormant. Fitness is frequency-dependent, as the success of each strategy depends on the composition of strategies within the local population. **(B)** Stress-induced payoff shift drives game switching. In a low-stress environment (left), the payoff structure resembles a Prisoner's Dilemma in which the proliferative strategy dominates (single Nash equilibrium at P vs. P). As stress increases (right), payoffs shift to a Snowdrift (Hawk–Dove) game structure in which neither strategy can exclude the other, generating stable coexistence through mixed Nash equilibria. This mechanism explains reversible intratumor heterogeneity without invoking genetic change. **(C)** Nested timescales of strategy dynamics. Tumor evolution operates across four coupled timescales, (i) fast metabolic reprogramming (Warburg ↔ OXPHOS switching; minutes to hours), (ii) medium-speed epigenetic remodeling (DTP state entry; days to weeks), (iii) slow genetic and chromosomal instability (CIN)-driven diversification (karyotype selection; weeks to months), and (iv) ultra-slow ecological succession with niche construction (r-selected → K-selected → cooperative; months to years). Faster dynamics operate within constraints set by slower processes, while slower dynamics are driven by cumulative changes at faster timescales.

Tumors as Populations Playing Repeated Games

Cancer evolution unfolds over repeated cycles of cell division, environmental fluctuation, immune engagement, and therapeutic intervention. Each cycle constitutes a new round of interaction in which the relative payoffs of strategies are recalculated. Unlike one-shot games, repeated games allow history, adaptation, and population structure to shape outcomes.

Evolutionary game models applied to cancer have demonstrated that interactions between cancer cells, stromal cells, and immune cells can stabilize cooperative and competitive behaviors. Examples include the production of shared growth factors, extracellular matrix remodeling, and immune suppression, all of which resemble public goods dynamics studied extensively in evolutionary biology[12,13]. In ovarian cancer, a subset of cells produces VEGF, benefiting all cells through enhanced angiogenesis, yet bearing the metabolic cost of secretion. Game-theoretic models predict that such producer cells should be susceptible to exploitation by non-producing "cheaters", yet producers persist across patient samples[36]. Experimental

tracking reveals that VEGF production is not constitutive but induced by hypoxia, and that spatial clustering protects producers from complete exploitation, as they disproportionately benefit from locally enhanced vascularization[37].

A recently discovered mechanism provides molecular evidence for how such cooperative equilibria are actively maintained. In breast cancer, miR-125b-high altruistic cells secrete trophic factors (IGFBP2 and CCL28) that confer chemotolerance to neighboring cells, while simultaneously activating a GAB1-PI3K-AKT-miR-125b lateral inhibition circuit that prevents those neighbors from becoming altruistic themselves[38]. This lateral inhibition enforces a stable sparse distribution of cooperators within the population, a molecular mechanism that directly implements the frequency-dependent equilibrium predicted by nonlinear public goods theory.

Spatial structure further enriches these dynamics. Local interactions can protect cooperative strategies from exploitation and allow multiple strategies to persist simultaneously. This aligns with empirical observations from spatial transcriptomics and intravital imaging, which reveal structured ecological niches within tumors rather than well-mixed populations[39]. Single-cell RNA sequencing combined with spatial transcriptomics in colorectal cancer has identified at least four metabolically distinct niches, including a proliferative niche near functional vasculature characterized by high MKI67 and glycolytic gene expression, a hypoxic-invasive niche at the tumor edge with EMT and HIF signatures, a nutrient-deprived core with autophagy and stress response markers, and an immune-interface niche enriched for interferon response genes[40]. Cells within each niche exhibit convergent transcriptional states despite diverse mutational backgrounds, indicating that micro-environmental context dominates over cell-intrinsic genotype in determining phenotype.

The repeated nature of the game also creates opportunities for conditional strategies that respond to history. Cells may "remember" prior stress exposure through epigenetic modifications, priming subsequent responses. Following transient hypoxia, breast cancer cells exhibit persistent up-regulation of HIF target genes even after reoxygenation, a phenomenon mediated by histone modifications at HIF-responsive promoters[41,42]. This memory allows cells to anticipate recurring stress, shifting payoffs in ways that static models cannot capture.

A critical consequence of repeated games played over tumor developmental time is ecological succession, *a change of seasons* in which the systematic, directional shift in strategy composition tracks the tumor's progressive modification of its own environment. Early-stage tumors, well-vascularized and nutrient-rich, favor rapidly proliferating, glycolytic strategies analogous to r-selected life histories: strategies that maximize intrinsic growth rate at the cost of competitive or stress-survival capacity[43,44]. As the tumor grows, however, its own metabolic activity generates hypoxia, acidosis, and nutrient depletion, while immune infiltration intensifies. These self-generated stresses shift selection toward stress-tolerant, metabolically flexible, and immune-evasive strategies, a K-selected regime, in which cells prioritize competitive efficiency and stress survival over proliferative speed[43,44]. In late-stage or heavily treated tumors, extreme environmental fragmentation favors cooperative strategies (senescence–SASP, metabolic symbiosis) and bet-hedging plasticity[45]. Crucially, this succession is not merely passive, as tumors actively engineer their microenvironment through angiogenesis, ECM remodeling, immune suppression, and metabolite secretion, progressively reshaping the payoff matrix to favor their current strategy portfolio, a form of niche construction in which the game's rules co-evolve with its players. Game-theoretic models that treat the payoff matrix as externally fixed will therefore systematically mispredict equilibrium states, because the matrix itself is a function of tumor developmental history and cumulative niche modification. This principle operates from the very earliest stages of tumorigenesis. In nascent squamous tumours of the mouse oesophagus, only $\sim$30% of incipient lesions develop a remodelled stromal scaffold, yet it is precisely these Niche+ tumours that persist long-term while Niche$-$ lesions are progressively eliminated[46]. The mechanism involves an EGF–SOX9-FN1 signalling axis in which a stress-activated epithelial subpopulation (SOX9-high, AREG-secreting) recruits PDGFRα-low lamina propria fibroblasts that deposit a fibronectin-rich scaffold beneath the tumour. A second epithelial subpopulation lacking SOX9 does not drive stromal remodelling yet benefits from the resulting niche, a producer–scrounger game at the founding stage. Tissue recombination experiments demonstrated that the tumour-derived stroma alone is sufficient to confer tumour properties to normal, carcinogen-free epithelium, showing that niche construction can rewrite the payoff matrix before the tumour has even established itself as a stable population.

Beyond the canonical cellular and stromal players, the nervous system represents an underappreciated participant in the tumor microenvironmental game. Tumors actively recruit neural innervation through neoneurogenesis, establishing a bidirectional tumor-nerve crosstalk in which cancer cells manipulate sympathetic and parasympathetic innervation to modulate stress responses, metabolic signaling, and immune function[6]. From a game-theoretic perspective, neural-tumor interactions introduce temporally varying signals, reflecting systemic neuroendocrine states and tissue-specific neural circuit properties, that dynamically reshape the payoff structures governing tumor cell strategies. Neurotransmitters and neurotrophic factors released by recruited neurons can promote proliferation, invasion, and therapeutic resistance, effectively adding an external player whose signaling inputs alter the relative fitness of competing stress-adaptive strategies within the tumor population. This neural dimension remains largely unexplored within formal evolutionary game models of cancer, yet its capacity to link systemic physiological states to local microenvironmental payoffs suggests it may substantially influence the dynamics of strategy selection and equilibrium formation.

Similarly, the intratumoral microbiota can be reframed as a dynamic player in the evolutionary game rather than a static environmental parameter. Microbial activities alter drug exposure, metabolic resource availability, and immune polarization within the tumor microenvironment. Within the repeated-game framework, these activities directly modify the payoff matrix. Bacterial metabolism of gemcitabine in pancreatic cancer, for instance, reduces the cost of proliferative strategies during chemotherapy by locally attenuating drug-induced stress[7], while *Fusobacterium nucleatum*-driven TLR4/MYD88 signaling in colorectal cancer enhances the fitness benefit of autophagy-dependent survival under therapeutic pressure[47]. Microbially derived metabolites further reshape the immune dimension of the game by polarizing myeloid and lymphoid populations, shifting the relative advantage of immunosuppressive versus immunosensitive tumor phenotypes[8,48]. Critically, the microbiome's own metabolic phenotype is itself subject to selection by the tumor microenvironment, creating a co-evolutionary dynamic in which microbial and tumor strategies jointly determine equilibrium outcomes. Incorporating the microbiome as a strategic agent in this sense represents a critical frontier for extending game-theoretic models of tumor evolution beyond cell-autonomous and stromal interactions.

Immune evasion as a multiplayer game

The interaction between tumors and the host immune system represents a prototypical multiplayer evolutionary game in which tumor cells and multiple immune cell populations, cytotoxic T lymphocytes, natural killer cells, macrophages, and regulatory T cells, compete for resources and engage in reciprocal strategy escalation within a spatially structured environment. Tumor cells employ a repertoire of frequency-dependent strategies to evade immune recognition and elimination, including down-regulation of major histocompatibility complex class I (MHC-I) molecules, expression of checkpoint ligands such as PD-L1, secretion of immunosuppressive cytokines (TGF-β and IL-10), and metabolic competition within the tumor microenvironment through glucose and amino acid sequestration[49,50]. These strategies do not operate in isolation but rather constitute coupled strategies in which the payoff to any single approach depends on the frequency of competing immune phenotypes and the density of other tumor subclones employing overlapping evasion mechanisms. Importantly, many immune evasion mechanisms exhibit features of nonlinear public goods games[13], where cells expressing high levels of PD-L1 or secreting TGF-β bear the biosynthetic cost of these investments, yet the immunosuppressive benefits accrue to the entire local tumor population in a manner that saturates with increasing cooperator frequency, creating a classic free-rider scenario in which subclones with low ligand expression benefit from suppression generated by their neighbors while achieving higher proliferation rates[51]. This structure generates complex evolutionary dynamics wherein equilibrium frequencies of evasion strategies depend critically on the relative payoffs to investment in immunosuppression versus other growth-promoting phenotypes, and on the degree to which spatial segregation permits local exploitation of immunity-suppressed microenvironments.

Immune cells, conversely, face distinct strategic trade-offs that shape their behavior within this game. Cytotoxic effector T cells obtain high payoffs from tumor cell killing but incur substantial metabolic costs and are subject to exhaustion, a cell-intrinsic dysfunction characterized by upregulation of inhibitory receptors (PD-1, TIM-3, LAG-3) and loss of effector function, that may be viewed as an adaptive strategy limiting terminal activation in persistently antigenic environments[52]. Macrophages within the tumor microenvironment exhibit phenotypic plasticity, polarizing toward either pro-inflammatory M1 states (characterized by IL-12 production and direct tumor cell killing) or immunosuppressive M2 states (characterized by IL-10, TGF-β, and pro-angiogenic factor secretion), with the frequency-dependent payoff structure favoring M2 dominance in established tumors where the suppressed immune context rewards investment in angiogenic and matrix-remodeling activities[53]. The multivalent nature of this game, involving simultaneous interactions between tumor cells, cytotoxic T cells, NK cells, regulatory T cells, and macrophages, precludes simple two-player analysis and instead requires consideration of how the strategies of one population reshape the payoff landscape for others, creating scenarios in which tumor-derived immunosuppression selects simultaneously for T cell exhaustion and macrophage M2 polarization, both of which further reward investment in tumor immune evasion.

The dynamics of immunoediting, classically characterized as progressing through elimination, equilibrium, and escape phases[49], map naturally onto the repeated game framework, with successive rounds of immune pressure iteratively selecting for tumor variants with enhanced evasion capacity while simultaneously sculpting the immune repertoire toward reduced diversity and increased dysfunction. The spatial architecture of immune infiltration, manifested as immune-inflamed, immune-excluded, or immune-desert phenotypes that can coexist within heterogeneous tumors, represents distinct game equilibria arising from differential investment by tumor subclones in checkpoint ligand expression, angiogenic factors, and physical barriers such as fibroblast-rich stroma that exclude immune cell infiltration[54]. The observation that tumor mutational burden (TMB) correlates positively with immunotherapy response in some cohorts but negatively in others illuminates the underlying game-theoretic complexity. High TMB generates neoantigen diversity that can sustain T cell recognition across multiple tumor variants, yet also increases the probability that subclones acquire immune evasion mutations in genes encoding MHC-I components, PTEN, or interferon-pathway effectors. When neoantigen clonal heterogeneity is high, selection favors spatially localized immune evasion strategies that permit subclone-specific escape while maintaining overall tumor growth[55]. The predictive value of TMB for immunotherapy response thus depends critically on whether the neoantigen repertoire is shared across the tumor population

or fragmented across spatially separated subclones, a distinction that reveals how raw mutation counts obscure the strategic interactions that determine tumor-immune equilibrium.

This perspective suggests that therapeutic interventions, whether checkpoint inhibitors, adoptive T cell therapies, or macrophage repolarization strategies, function by fundamentally altering the payoff structure of the tumor-immune game, shifting equilibrium frequencies of immune activation and tumor evasion, and thereby creating transient windows of competitive imbalance during which immune-mediated tumor control becomes possible. Critically, how tumor cells die under therapeutic pressure is itself a strategic variable. Immunogenic cell death (through ferroptosis, necroptosis, or pyroptosis) releases damage-associated molecular patterns that activate anti-tumor immunity, effectively penalizing surviving neighbors, whereas non-immunogenic death (apoptosis) permits silent clearance that leaves the remaining population's payoff structure intact[56]. The modality of cell death thus determines whether therapy strengthens or weakens the immune dimension of the game for survivors. Understanding these interactions as frequency-dependent evolutionary games rather than as linear cause-and-effect relationships provides both mechanistic insight into why immunotherapies fail in specific contexts and a framework for predicting which tumor-immune configurations are most vulnerable to strategic intervention.

Senescence as a sacrificial cooperative strategy

Genotoxic, oxidative, and proteotoxic stresses do not invariably eliminate the cells they damage; a substantial fraction instead enters therapy-induced senescence, converting acute stress into a chronic microenvironmental signal. The therapeutic induction of cellular senescence, historically regarded as a tumor-suppressive outcome of chemotherapy and radiation[57], thus represents a profound example of how stress adaptation itself becomes a cooperative strategy within the neoplastic ecosystem. Rather than constituting an unambiguous victory against cancer, therapy-induced senescent cells paradoxically transform their own stress-arrested state into a paracrine resource that may accelerate relapse and therapeutic resistance. These permanently cell-cycle-arrested cells, though incapable of direct proliferation, do not remain inert within the tumor microenvironment; instead, they secrete the senescence-associated secretory phenotype (SASP)[58], a constellation of cytokines, chemokines, growth factors, and extracellular vesicles that profoundly reshape the stress landscape and fitness of neighboring, non-senescent clonal siblings. This phenomenon exemplifies a "sacrificial public goods" game[13], fundamentally distinct from previously discussed cooperation mechanisms such as VEGF-mediated angiogenesis. Here, senescent cells bear a terminal, irreversible fitness cost, permanent cell cycle arrest and inability to contribute to future clonal expansion, while simultaneously generating substantial paracrine benefits for their non-senescent neighbors. The senescent cell subsidizes a collective good through extracellular vesicle (EV) secretion; recent evidence demonstrates that therapy-induced senescent colorectal cancer cells release SASP-enriched extracellular vesicles, including those enriched in the serine protease inhibitor SERPINE1, which are internalized by recipient cancer cells and promote their invasion, migration, and acquisition of drug-resistant phenotypes.[59] Through the lens of game theory, this cooperation strategy generates a frequency-dependent payoff in which the senescent cell sacrifices future reproduction but increases the immediate survival and proliferative capacity of non-senescent cohorts, thereby indirectly ensuring clonal persistence even as the senescent individual itself is eliminated. The SASP-mediated cooperative network exhibits the hallmark features of an iterated game played across multiple rounds within the microenvironmental neighborhood: senescent cells produce signals (public goods), non-senescent cells benefit from elevated growth, invasion, and resistance, and the total fitness of the cooperating cohort exceeds what either strategy could achieve in isolation. Even the cells that therapy eliminates do not exit the game silently; they leave *tears in heaven*—SASP factors, extracellular vesicles, and damage signals that persist in the microenvironment and reshape the adaptive landscape for those that survive.

This duality exposes a fundamental therapeutic paradox, as regimens that successfully trigger senescence may inadvertently enrich a cooperating subset that drives relapse.[60] The paradox emerges naturally from the game-theoretic framework because the payoff matrix governing senescence shifts with therapeutic context. During drug exposure, senescence represents a favorable equilibrium in which cells that surrender proliferative capacity avoid the costs of division in a toxic microenvironment. Upon therapy withdrawal, the payoffs restructure, as senescent cells cannot participate in clonal expansion, yet their SASP output creates a niche of elevated growth factor signaling and enhanced invasion cues that disproportionately benefit non-senescent subclones. The game transforms from "sacrifice for survival" into "exploitation of altruists," where the senescent cell's evolutionary contribution is decoupled from its own fitness. Senescence persists as an evolutionarily stable strategy not because senescent cells themselves proliferate, but because their presence elevates population-level fitness, ensuring propagation of the genetic and epigenetic machinery that enables senescence induction.

Metabolic symbiosis as a cooperative game

Metabolic and nutrient stress, arising from poor vascularization, glucose depletion, and oxygen limitation, is among the most pervasive selective pressures in solid tumors. Rather than each cell independently solving this metabolic crisis, tumors frequently evolve an inter-cellular division of metabolic labor that distributes the burden of stress adaptation across cell types. The metabolic partnership between cancer cells and cancer-associated fibroblasts (CAFs), encapsulated in the concept of the "reverse Warburg effect," exemplifies a two-player cooperative game in which asymmetric payoffs nevertheless maintain

stable cooperation through mutual fitness benefits under shared metabolic constraint. In this metabolic symbiosis, CAFs undergo a dramatic metabolic reprogramming orchestrated by paracrine signals from oncogenic tumor cells, shifting their metabolism from oxidative phosphorylation to aerobic glycolysis, the classical Warburg effect normally associated with cancer cells themselves. These reprogrammed stromal cells then export their glycolytic products (lactate, pyruvate, and ketone bodies) across the tumor-stromal interface via monocarboxylate transporters (MCT4 in CAFs, MCT1 in cancer cells), providing high-energy substrates that fuel oxidative phosphorylation within cancer cells.[61] In game-theoretic terms, this interaction differs from the symmetric public goods games typically discussed in the cooperation literature[13]. Rather than equivalent players choosing whether to contribute to a shared resource, the reverse Warburg effect represents an asymmetric coercive game in which one player (the tumor cell) enforces metabolic cooperation from another (the CAF). The resulting Nash equilibrium is maintained through repeated interaction, wherein CAFs incur a fitness cost through the energy and biosynthetic resources devoted to metabolite export, while cancer cells gain the energetic advantages of oxidative metabolism without the oxidative stress that would accompany their own glycolytic-to-oxidative transition. Crucially, tumor cells actively weaponize oxidative stress to enforce this metabolic division of labor, generating reactive oxygen species that both drive CAF reprogramming and create a selective pressure favoring cancer cells with increased capacity for oxidative phosphorylation.

This coupling reveals an essential asymmetry, as the tumor cell functions as a metabolic parasite, extracting subsidized fuel while imposing glycolytic labor on its stromal partners through paracrine ROS signaling that drives CAF reprogramming through selective pressure rather than reciprocal negotiation.[62] CAFs persist in this exploitative arrangement not because it favors their fitness in absolute terms, but because resistance to tumor-derived signals invokes even costlier selective penalties, a stable equilibrium in an asymmetric game where the parasite sets the rules. Therapeutically, this asymmetry creates a specific vulnerability, as disrupting MCT-mediated lactate and pyruvate transport collapses the metabolic payoff matrix, forcing cancer cells to either revert to glycolytic metabolism (with its attendant oxidative stress) or face energetic insufficiency.[61] Because the tumor-CAF relationship is a directed parasitism rather than a symmetric mutualism, therapies that dismantle it should encounter reduced evolutionary resistance compared to those targeting symmetric interactions. This metabolic symbiosis illustrates how stress adaptation in tumors extends beyond cell-autonomous responses, encompassing the collective management of metabolic stress through inter-cellular coupling represents a population-level adaptive strategy whose stability and vulnerability are best understood through the game-theoretic lens.

Whereas tumor–CAF metabolic coupling represents an asymmetric, coercive game, recent work has uncovered a symmetric cooperative game played among tumor cells themselves under nutrient stress. Guzelsoy and colleagues[63] demonstrated that amino acid-starved tumor cells collectively secrete the aminopeptidase CNDP2, which digests extracellular oligopeptides into free amino acids that function as a classic public good, wherein secreting cells bear the biosynthetic cost while all neighboring cells, including non-secretors, benefit from the liberated nutrients. This cooperative nutrient scavenging exhibits a pronounced Allee effect, in which population fitness increases with cell density because the shared aminopeptidase pool scales with the number of contributors, creating a positive feedback loop that sustains proliferation under conditions where isolated cells would starve. Critically, disruption of CNDP2 activity, either genetically or through the aminopeptidase inhibitor bestatin, drove tumor populations toward a critical extinction threshold both in vitro and in vivo, demonstrating that cooperation itself constitutes a targetable vulnerability. This finding provides direct experimental evidence that the public goods framework, previously applied to growth factor secretion and immune suppression, extends to the most fundamental level of tumor metabolism, namely the collective acquisition of amino acids required for biosynthesis and survival.

Plasticity as Strategy Switching

Plasticity beyond noise

Phenotypic plasticity is among the most reproducible and quantitatively documented features of cancer. Single-cell RNA sequencing of melanoma tumors revealed that individual tumors harbor cells spanning a continuum between a MITF-high proliferative state and an AXL-high invasive state, with these two programs present at varying proportions across patients yet drawn from the same transcriptional repertoire[25]. Analogous discrete-yet-reversible cell state architectures have since been identified across cancer types, including proliferative, mesenchymal, neural crest-like, and drug-tolerant states in melanoma[64]; classical and basal-like programs in pancreatic adenocarcinoma; and neuroendocrine versus luminal states in prostate cancer[65]. The consistency of these architectures across genetically distinct tumors argues against stochastic noise. However, these recurrent cell states are not strategies in themselves, they are the transcriptional signatures from which strategies must be inferred by considering the game state each cell faces. A cell in a "mesenchymal" transcriptional state may be executing an invasive strategy or a drug-tolerant strategy depending on the local configuration of stresses and neighbors. We propose that tumor heterogeneity should be understood as the structured outcome of frequency-dependent selection shaped by recurring selective pressures, not as a barrier to treatment but as a predictable feature amenable to game-theoretic analysis.

In fluctuating environments, evolutionary theory predicts that conditional strategies outperform fixed ones. Plasticity allows organisms to match phenotype to context, even at the cost of short-term inefficiency, a principle validated across biological

systems from bacterial persistence to developmental canalization.

The emergence of drug-tolerant persister (DTP) cells provides a clear example[21,66]. In EGFR-mutant lung cancer treated with EGFR inhibitors, a rare subpopulation survives initial treatment without acquiring resistance mutations. These DTP cells exhibit slow proliferation, elevated oxidative phosphorylation, and IGF1R pathway activation. Critically, DTP cells are not genetically distinct from sensitive cells; single-cell lineage tracing demonstrates that sensitive cells can stochastically enter the DTP state, and DTP cells can revert to sensitivity upon drug withdrawal. The transition is mediated by chromatin remodeling at enhancers controlling IGF1R and other receptor tyrosine kinases, allowing dynamic access to alternative signaling strategies. Similar phenomena have been observed across cancer types and therapeutic modalities, including platinum-induced DTPs in ovarian cancer, MEK inhibitor-induced DTPs in melanoma, and androgen deprivation-induced neuroendocrine plasticity in prostate cancer[65].

Strategy switching in evolutionary games

Within a game-theoretic framework, plasticity enables strategy switching in response to changing game states and their associated payoffs. Cancer cells need not commit permanently to a single strategy; instead, they maintain epigenetic access to multiple cell states, and therefore to the strategies those states can implement, transitioning between them as the stress landscape and neighbor composition change. The rapid emergence of drug-tolerant persister states and their disappearance upon drug withdrawal exemplifies this principle: the game state shifts when therapy is applied (altering the payoff structure), cells transition to a persister cell state that implements a tolerance strategy, and upon drug withdrawal the game state reverts, favoring transition back to a proliferative state and its associated strategy.

However, not all switching is equivalent, and different switching modes enter the game with distinct evolutionary consequences. Three modes should be distinguished. *Stochastic (intrinsic) switching* occurs at a basal rate independent of environmental cues and enters the game as a mixed-strategy parameter; the switching rate itself is evolvable, and its optimum depends on the frequency and amplitude of environmental fluctuations. Drug-tolerant persister entry at baseline ($\sim 10^{-5}$–10^{-4} per cell per generation) exemplifies this mode. *Environmentally induced (conditional) switching* is payoff-responsive, meaning that cells detect stress cues and transition deterministically or semi-deterministically to an alternative state, entering the game as a best-response function. Hypoxia-induced HIF program activation and therapy-induced persister entry are canonical examples. *Irreversible diversification* expands the strategy set itself rather than redistributing cells among existing strategies; CIN-driven karyotypic changes and loss-of-function mutations in immune evasion genes exemplify this mode. Because stochastic switching preserves reversibility, it sustains mixed-strategy equilibria; conditional switching enables rapid equilibrium tracking as payoffs shift; and irreversible diversification creates new strategic types that may destabilize existing equilibria entirely. Therapeutic implications differ accordingly, as plasticity-restricting therapies (e.g., epigenetic inhibitors) target the first two modes but are irrelevant to the third, while agents that reduce chromosomal instability constrain the third without affecting epigenetic switching.

Formal models of phenotypic switching in cancer have demonstrated that bet-hedging strategies, in which cells stochastically transition between states at low frequencies, can be evolutionarily stable in fluctuating environments where no single phenotype dominates across all conditions. The optimal switching rate depends on the timescale and amplitude of environmental fluctuations. When stress varies rapidly relative to cell division time, high switching rates are favored; when stress is persistent, low switching rates with strong environmental responsiveness dominate. Empirical measurements of switching kinetics in melanoma cells exposed to BRAF inhibitors reveal switching rates consistent with the timescale of clinical drug holidays, suggesting that plasticity dynamics are tuned to the predictable rhythms of treatment.

This perspective reconciles stochastic and regulated views of plasticity by recognizing variability itself as an evolved property of systems under chronic stress. Plasticity is neither purely deterministic nor purely random but reflects a controlled stochastic process, analogous to bacterial persistence mechanisms that generate phenotypic heterogeneity through noisy gene expression within a regulated framework.

Yet if plasticity is so advantageous, why do not all tumors maximize it? The answer is that plasticity itself is costly, maintaining open chromatin at alternative-state enhancers, sustaining the signaling machinery for state detection, and tolerating the risk of maladaptive switching all impose fitness penalties. Plasticity is therefore subject to its own frequency-dependent selection. Comparative studies across taxa demonstrate that the evolution of phenotypic plasticity depends critically on environmental predictability, such that plasticity is favored when environmental fluctuations are rapid relative to generation time, large in magnitude, and accompanied by reliable cues that predict future conditions[67]. When stress is constant and predictable, as in chronically hypoxic tumor cores, fixed, specialized strategies outperform plastic ones because they avoid the overhead of maintaining epigenetic flexibility. Conversely, in regions experiencing cycling hypoxia, intermittent immune engagement, or scheduled drug holidays, the payoff to plasticity increases because no single fixed strategy dominates across all conditions. This logic explains why drug-tolerant persister states are rare at baseline ($\sim 10^{-5}$–10^{-4} per cell per generation), indicating that the population occupies a mixed-strategy equilibrium in which most cells remain in the high-fitness proliferative

state while a small fraction bet-hedges in the persister state, with observed switching rates consistent with theoretical optima given typical fluctuation timescales[21]. The prediction follows that plasticity rates should be highest in tumors experiencing unpredictable or rapidly alternating stresses and lowest in tumors under stable selective pressure, a prediction directly testable through longitudinal single-cell tracking under controlled fluctuating versus constant stress regimes. A striking mechanistic example is sublethal or "minority" mitochondrial outer membrane permeabilization (MOMP), in which only a fraction of mitochondria within a cell undergo apoptotic pore formation, releasing limited cytochrome c that activates caspases at sublethal levels[56]. Rather than triggering death, this partial apoptotic signaling engages caspase-activated DNase to generate low-level DNA damage, fueling mutagenesis and karyotypic diversification without eliminating the cell. Sublethal MOMP thus functions as an endogenous bet-hedging mechanism that converts incomplete stress responses into phenotypic exploration: cells that narrowly survive apoptotic signaling emerge with expanded mutational repertoires and altered epigenetic states, effectively sampling new strategic positions at a cost well below that of full CIN.

The adaptive dynamics of tumors operate simultaneously across multiple nested temporal timescales, each governing distinct categories of strategy switching and each imposing constraints upon the others in a manner directly analogous to nested game theory as developed in political science and economic governance. Fast timescale games, operative on the order of minutes to hours, involve acute metabolic decision-making and rapid epigenetic responses to transient microenvironmental perturbations, metabolic switching between glycolysis and oxidative phosphorylation, rapid engagement or disengagement of hypoxia response pathways, immediate modulation of growth factor receptor signaling in response to paracrine cues. Medium timescale games, occurring across days to weeks, involve the sustained epigenetic plasticity previously discussed, wherein histone modifications and DNA methylation patterns are rewritten to support stable transitions between phenotypic states while maintaining genetic continuity. Slow timescale games, unfolding across weeks to months, involve the genetic mutations and chromosomal rearrangements that are fundamentally irreversible at the timescale of individual cells but constitute the heritable substrate upon which all faster games are played. This temporal nesting creates a hierarchical constraint structure, a *train of thought* in which each timescale's logic constrains the next, such that strategies available at the fast timescale are determined by the epigenetic state established at the medium timescale, which in turn is constrained by the genetic architecture created through slow-timescale mutations (**Fig. 2C**). Critically, therapeutic interventions also operate at specific timescales, and the efficacy of a given treatment is intimately related to whether its timescale of action matches or mismatches the timescale of the dominant adaptive game being played. A chemotherapeutic agent that acts on the fast timescale (acute metabolic disruption) may be rapidly overcome through medium-timescale epigenetic plasticity if tumor populations have evolved sufficient bet-hedging capacity. Conversely, a therapy targeting slow-timescale genetic mutations (such as CRISPR-based approaches) may be circumvented through rapid fast-timescale metabolic or medium-timescale epigenetic compensation. The nested-games framework predicts that the most durable therapeutic strategies will be those that simultaneously target adaptive mechanisms across all three timescales, preventing rapid compensation at any single level. This theoretical prediction carries immediate experimental and clinical implications, because combination therapies that coordinately attack metabolic plasticity (fast), epigenetic plasticity (medium), and genetic instability (slow) should encounter substantially reduced adaptive escape compared to monotherapies targeting single timescales.

Chromosomal instability as a generator of strategic diversity

Chromosomal instability (CIN), the ongoing acquisition of whole-chromosome gains and losses, along with larger-scale structural rearrangements, has long been viewed principally as a hallmark of malignancy and a driver of genomic disorder. However, reframing CIN through the lens of game-theoretic strategy switching reveals it as a potent mechanism for generating karyotypic heterogeneity that expands the phenotypic space available to tumor populations. Rather than pathological noise, CIN functions as *systematic chaos*—a bet-hedging mechanism of extraordinary scope, because by continuously generating diverse aneuploid karyotypes within a single population, CIN-high tumors maintain an expansive portfolio of potential adaptive strategies, each embodied in distinct chromosomal compositions. This karyotypic heterogeneity directly translates into strategic diversity because whole-chromosome gains and losses alter the dosage of large gene cohorts simultaneously, particularly stress-response regulators, metabolic enzymes, and immune evasion factors, creating an array of phenotypic states that would be inaccessible through point mutations or epigenetic modulation alone[68,69]. The population thus explores an enlarged strategy space, sampling multiple biochemical regimes and phenotypic configurations in parallel, a form of distributed search across the landscape of possible adaptive solutions.

The fitness cost of this exploration is substantial, with the majority of aneuploid karyotypes being inviable or exhibiting reduced proliferative capacity relative to euploid progenitors[68]. Centrosome amplification, a structural driver of CIN, similarly imposes metabolic and redox burdens that generate exploitable vulnerabilities[70]. Yet CIN also generates outward-facing strategic benefits, as micronuclei from chromosome missegregation activate cGAS–STING signaling, paradoxically promoting invasion through NF-κB rather than triggering immune clearance[68], while the oxidative stress from centrosome amplification drives an extra centrosomes-associated secretory phenotype (ECASP) that promotes non-cell-autonomous invasion[71]. The CIN trade-off

is thus double-edged, with inward-facing metabolic costs constraining the CIN-high cell while outward-facing paracrine signaling expands invasive capacity for the local population, a structure reminiscent of the senescence–SASP cooperative game. This trade-off is environment-dependent. Under stable stress, CIN is wasteful because most variant karyotypes are less fit than the current optimum. However, when stress fluctuates rapidly or takes novel forms, CIN provides raw material for rapid adaptive escape, as pre-sampled karyotypic diversity means the population likely harbors variants closer to the new optimum. This mirrors the bet-hedging logic invoked for epigenetic plasticity but operates at a different timescale, karyotypic changes are largely irreversible and heritable, and through a different mechanistic substrate, namely gene dosage rather than promoter accessibility.

Empirical evidence supports this strategic interpretation. CIN burden correlates with therapy resistance and poor prognosis across solid tumors, consistent with more rapid access to drug-resistant phenotypes[68,72], and experimental induction of CIN in previously stable cell lines accelerates adaptation to chemotherapy[69]. Yet CIN-high tumors often exhibit enhanced responses to checkpoint immunotherapy, attributable to the increased neoantigen burden that accompanies karyotypic instability[72]. This divergence, rapid adaptation to chemical stress but vulnerability to immune pressure, illustrates that CIN expands the accessible strategy set in directions that are simultaneously advantageous and costly. This framework generates a specific prediction. CIN rates should be highest in tumors experiencing rapidly fluctuating or unpredictable stress environments (e.g., cycling hypoxia, intermittent immune pressure), where the exploratory benefit of karyotypic diversity outweighs its fitness cost, and lowest in tumors under stable selective pressure, where exploitation of the current optimum dominates.

Metastasis as Strategy Export and Game Migration

The preceding sections have examined how stress adaptation operates within the primary tumor, where cells adopt, switch, and cooperate around strategies shaped by stress. However, cancer is not confined to its organ of origin. The process of metastatic dissemination introduces a fundamentally novel dimension to the stress adaptation framework, namely the spatial migration of stress-evolved strategies to entirely novel ecosystems where the stress landscape, and therefore the payoff matrices, are not merely quantitatively different but qualitatively distinct. Disseminated tumor cells (DTCs) represent populations of cancer cells that have successfully exported their evolved strategies, metabolic plasticity, immune evasion, stress resistance, and cooperative phenotypes, to foreign tissue microenvironments, including bone marrow, hepatic parenchyma, and cerebral vasculature. Within each of these novel organs, the evolutionary landscape that selected for specific strategies in the primary tumor becomes not merely suboptimal but potentially maladaptive. A strategy that proved evolutionarily stable at high frequency within the permissive, well-vascularized primary tumor microenvironment, where the dominant stresses were nutrient competition and immune pressure, may impose severe fitness costs when introduced into the bone marrow niche, where hypoxia is profound, oxidative stress is distinct, immune surveillance operates through different effector populations, and the cellular composition of the microenvironment is fundamentally alien. The metastatic process thus represents a game-theoretic "reset": cells carrying strategies optimized for one set of payoffs must rapidly re-equilibrate to an entirely reorganized payoff matrix (**Fig. 3A**). This reframing of metastasis as strategy export and migration across disparate microenvironmental games opens a novel theoretical frontier, one in which dormancy, awakening, and organotropic specificity acquire explicit game-theoretic interpretations.

Before reaching foreign organs, however, disseminating cells must survive the circulatory compartment, a distinct game environment imposing stresses absent from the primary tumor, including anoikis from loss of ECM attachment, shear forces ranging from 1–30 Dyn cm^{-2} in arterial flow, and direct immune surveillance by NK cells[56]. Circulating tumor cell (CTC) clustering represents a cooperative strategy for navigating this hostile compartment. Homotypic clusters maintain cell–cell contacts through plakoglobin, E-cadherin, and CD44, reducing intracellular ROS, inhibiting ferroptosis, and shielding interior cells from NK cell-mediated killing[56]. Heterotypic clusters incorporating platelets, neutrophils, or cancer-associated fibroblasts extend this cooperative network further, with platelets conferring anoikis resistance through YAP1 activation and physical shielding from immune recognition. From a game-theoretic perspective, CTC clustering constitutes a cooperative game with clear frequency-dependent costs (larger clusters face increased capillary entrapment) and benefits (dramatically enhanced survival), whose equilibrium cluster size reflects the local balance of circulatory stresses. The successful colonizers at distant sites are therefore not random emigrants but survivors of a distinct intermediate game played under rules entirely different from both the primary tumor and the eventual metastatic niche.

Moreover, each target organ presents a qualitatively different payoff matrix, requiring organ-specific metabolic rewiring. Lung DTCs activate pyruvate carboxylase-dependent anaplerosis to exploit the pyruvate-rich lung interstitium, while brain metastases shift to branched-chain amino acid catabolism and even form pseudo-tripartite synapses with glutamatergic neurons to access glutamate[56]. These organ-specific metabolic adaptations represent distinct strategy sets that are inaccessible to cells lacking sufficient plasticity, providing a mechanistic basis for the observation that metastatic organotropism reflects not merely adhesion molecule compatibility but metabolic strategy compatibility with the foreign microenvironment's payoff structure.

The dormant phenotype of disseminated tumor cells in foreign microenvironments can be understood as a stress adaptation strategy, a "wait-and-switch" response to the novel and inhospitable stress landscape of the foreign organ. Dormancy is not a

passive, energetically neutral state but an active, metabolically constrained equilibrium, analogous to the quiescent stress-tolerant states discussed in the context of drug-tolerant persisters, in which cells maintain minimal biosynthetic activity and DNA replication while simultaneously scanning the microenvironmental landscape for signals indicating a shift in the stress regime. From an evolutionary game theory perspective, this represents a best-response strategy when the cost of proliferation (exhaustion of niche resources, immune recognition, accumulation of deleterious mutations) exceeds the benefit. The fundamental question is under what conditions the payoffs shift such that dormancy becomes disadvantageous and proliferative growth becomes the evolutionarily stable strategy. The answer lies in microenvironmental remodeling, a process initiated not by the DTCs themselves but by molecular scouts that precede the physical arrival of metastatic cells. Tumor-derived exosomes bearing specific integrin signatures engage with resident cells of the distant organ, progressively reshaping the immune, metabolic, and stromal landscape to resemble the permissive microenvironment of the primary tumor. These EV-derived advance signals function as a form of "microenvironmental engineering", effectively pre-conditioning the foreign tissue such that upon the DTC's arrival, the payoff matrix has been partially remodeled to favor growth rather than dormancy.[73] Hoshino and colleagues demonstrated that tumor-derived exosomes bearing specific integrin signatures ($\alpha 6\beta 4$ and $\alpha 6\beta 1$ for lung, $\alpha v\beta 5$ for liver) selectively accumulate in their cognate target organs and interact with resident cells to form a permissive pre-metastatic niche, thereby establishing an organotropic specificity at the molecular level. This exosome-mediated pre-conditioning can be conceptualized within game theory as an ex ante manipulation of the payoff matrix, the DTCs do not arrive to an unchanged game but to a game that has been strategically altered by their own metastasis-promoting signals.

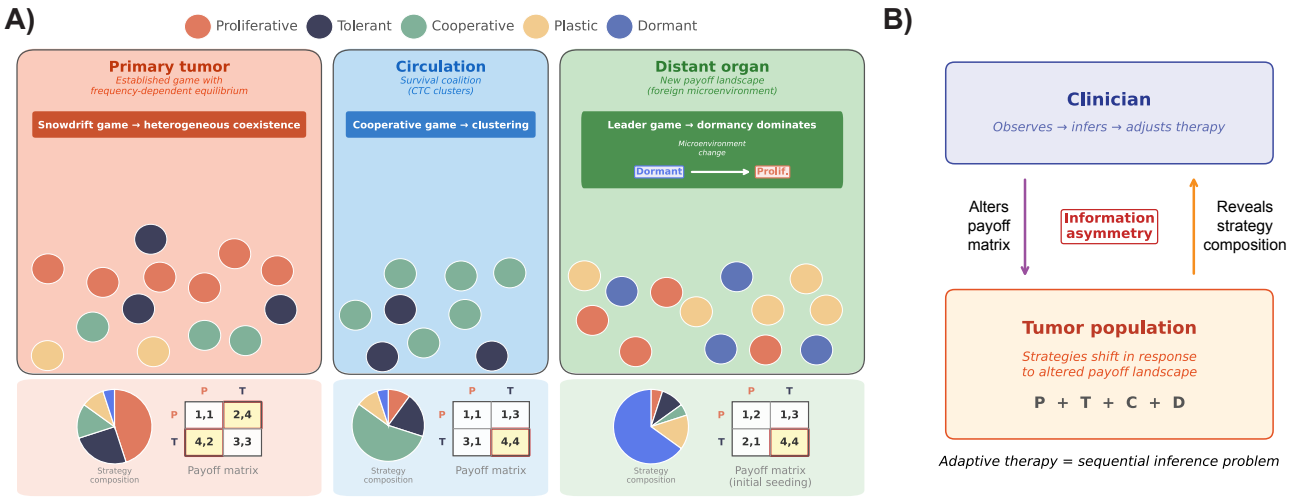

Figure 3. Metastasis as game migration and therapy as metagame. (A) Metastasis as migration between games with distinct payoff landscapes. Each anatomical compartment, primary tumor, circulation, and distant organ, imposes a different payoff matrix on the tumor population, selecting for different equilibrium strategy compositions. In the primary tumor, a Snowdrift game maintains heterogeneous coexistence of proliferative and tolerant strategies. During circulation, a cooperative game favors CTC clustering. At the distant organ, a Leader game initially selects for dormancy, with subsequent microenvironmental changes triggering a game switch to proliferative reactivation. Strategy composition pie charts and mini payoff matrices illustrate how the dominant game type and equilibrium shift across compartments. **(B)** Therapy as a Stackelberg metagame. The clinician (leader) and tumor population (follower) engage in a sequential game characterized by information asymmetry. The clinician observes tumor behavior, infers strategy composition, and adjusts therapy to alter the payoff matrix; the tumor population responds by shifting strategies toward the new equilibrium imposed by the altered fitness landscape. Adaptive therapy exploits this structure as a sequential inference problem, modulating treatment to maintain a favorable competitive balance rather than seeking maximum cell kill.

The transition from dormancy to proliferation at the metastatic site represents a game-theoretic phase transition analogous to critical phenomena in physics. During the dormancy phase, the equilibrium payoff to the DTC favors quiescence; the strategy "remain metabolically minimal, maintain chromosomal stability, exploit stromal resources minimally" yields higher expected fitness than "proliferate aggressively." However, as microenvironmental remodeling accumulates through exosome-mediated signaling and the recruitment of immunosuppressive and pro-angiogenic cells, a critical threshold is crossed and the dormant population is suddenly *awake*, the payoff matrix undergoes a discontinuous restructuring such that proliferation becomes the

dominant strategy. This phase transition is not driven primarily by cell-intrinsic changes but by the exogenous remodeling of the game itself. The temporal kinetics of the pre-metastatic niche formation thus directly determine the latency period between primary tumor dissemination and overt metastatic outgrowth. Microenvironmental studies using three-dimensional co-culture systems and organotypic models have demonstrated that the dormancy-to-growth transition is profoundly regulated by contextual cues; dormant breast cancer cells maintained in bone marrow-derived stromal co-cultures display quiescence-associated transcriptional signatures, while the same cells in a more permissive collagen-rich microenvironment rapidly shift to proliferative phenotypes.[74] This experimental evidence confirms the game-theoretic prediction that dormancy and proliferation are context-dependent strategies rather than fixed cell-intrinsic states. Furthermore, mathematical models of dormancy incorporating fitness-dependent payoffs demonstrate that DTCs maintain a form of bet-hedging within the dormancy strategy itself, a subpopulation remains in cycle, prepared to exploit local opportunities as soon as the microenvironmental game shifts.[75] This internal heterogeneity within the dormant population represents a mixed strategy equilibrium, wherein the population as a whole hedges against uncertainty regarding the timing and extent of microenvironmental permissiveness.

The game-theoretic reconceptualization of metastasis as strategy export across spatially heterogeneous microenvironmental games also illuminates the phenomenon of metastatic colonization failure. Most disseminated tumor cells do not successfully establish macrometastases; they are *falling into infinity*—scattered across vast, inhospitable tissue compartments where they undergo either apoptosis or indefinite dormancy, a phenomenon that remains mechanistically mysterious within conventional oncology frameworks. However, when viewed through the lens of game migration, colonization failure emerges as the consequence of maladaptation, DTCs carrying strategies optimized for the primary tumor microenvironment find themselves in a novel game where those strategies yield negative payoffs and no rapid re-equilibration is possible. The cells that do successfully colonize distant organs are not necessarily those with the most aggressive proliferative capacity but rather those carrying strategies compatible with the foreign microenvironment's payoff matrix. This predicts that the most frequent survivors of microenvironmental migration are those DTCs whose intrinsic plasticity allows rapid strategy switching, precisely the bet-hedging and epigenetic plasticity mechanisms discussed earlier in the framework. The metastatic paradox, that metastatic outgrowth requires cells with heightened stress resistance and plasticity, even though these traits reduce proliferative fitness in the primary tumor, resolves naturally when one recognizes that metastasis involves migration across games with incompatible payoff structures. The evolutionary bottleneck of metastasis thus selects for cellular and populational strategies optimized not for immediate fitness in the primary tumor but for strategic flexibility across heterogeneous stress landscapes. In this sense, metastasis represents the ultimate test of the stress adaptation framework, demanding not mastery of any single stress but the capacity to re-adapt when the entire stress environment changes.

Extracellular vesicles as long-range signaling strategies

A recurring theme throughout this framework is that stress adaptation is not solely a cell-autonomous process. Cells under stress broadcast signals that reshape the adaptive landscape for others. Extracellular vesicles (EVs), including exosomes, microvesicles, and apoptotic bodies, represent the most versatile and far-reaching medium for this inter-cellular stress communication, yet their role has been incompletely integrated into game-theoretic frameworks. Rather than representing mere "molecular garbage" or passive byproducts of cellular metabolism, EVs function as packages of stress-adaptive information, carrying stress-response transcripts, metabolic enzymes, and signaling molecules, that encode the sender's stress state and actively reshape the stress landscape and payoff structures for receivers. In the context of therapy-induced senescence, SASP-enriched extracellular vesicles secreted by senescent cells constitute a form of "weaponized cooperation", senescent cells package growth-promoting, invasion-supporting, and drug-resistance-conferring molecules into stable, transferable particles that can persist in the microenvironment and be internalized by non-senescent recipients even after the senescent cell itself has been cleared by immune mechanisms. This allows the senescent cell's cooperative strategy to achieve fitness effects that outlast the senescent cell itself, effectively extending its evolutionary influence beyond its biological lifespan. In the context of metastatic dissemination, tumor-derived exosomes bearing specific integrin signatures constitute an even more sophisticated strategic deployment, as these EVs are not merely signals of intent but active remodelers of the distant microenvironment, selectively accumulating in target organs and engaging resident cells to construct a pre-metastatic niche optimized for incoming DTCs. The exosome thus functions as a molecular scout or advance party, establishing beachheads in foreign tissue before the metastatic population itself arrives. Game-theoretically, EVs represent a means of extending the spatial and temporal reach of individual cells' stress-adaptive decisions, allowing coordination and payoff manipulation across scales that would otherwise exceed cellular communication capacity. By enabling stress-adapted cells to transmit their adaptive programs to naive recipients, EVs transform stress adaptation from a local, cell-intrinsic process into a distributed, population-level phenomenon, one in which the stress history of individual cells propagates through the tumor ecosystem and shapes the evolutionary trajectory of the entire population.

Therapy as a Perturbation of the Game

Most anticancer therapies are designed to kill cells, but their dominant evolutionary effect is often to reshape the stress landscape. Chemotherapy intensifies genotoxic and proteotoxic stress, targeted therapies remove specific signaling strategies, and immunotherapies amplify immune-mediated pressure. Each intervention alters the payoff matrix governing which strategies are viable.

The game-theoretic perspective predicts that therapies do not merely select for resistance but fundamentally change the type of ecological interaction among tumor cell populations. This prediction was recently validated empirically by Kaznatcheev and colleagues[24], who directly measured the evolutionary games played by non-small cell lung cancer under different treatment conditions. They demonstrated that the addition of alectinib or removal of cancer-associated fibroblasts qualitatively shifted the game from a Leader game to a Deadlock game, confirming that interventions can transform not just the magnitude but the fundamental nature of cellular interactions. This provides experimental proof of concept for the principle of treating the game, not the player (**Fig. 3B**).

Evolutionary game models of cancer treatment emphasize that resistance does not require new mutations to arise; it often reflects the expansion of pre-existing or inducible strategies that were previously suppressed. This insight has motivated adaptive therapy approaches that explicitly aim to manage, rather than eliminate, resistant populations. In metastatic castration-resistant prostate cancer, adaptive therapy trials[76] have demonstrated that maintaining a population of treatment-sensitive cells can competitively suppress resistant cells during drug holidays, extending time to progression compared to continuous dosing. The underlying logic is game-theoretic, because in the absence of treatment, sensitive cells outcompete resistant cells, which typically bear fitness costs associated with resistance mechanisms. Therapy inverts this relationship, favoring resistant cells. By cycling treatment on and off, adaptive therapy prevents either population from reaching fixation.

More recent work has framed cancer treatment as a leader–follower problem, in which clinicians act as rational leaders anticipating the adaptive responses of tumor cell populations. While still largely theoretical, this approach underscores the importance of understanding therapy as a dynamic interaction rather than a static intervention. Optimal control models suggest that pulsatile dosing schedules, deliberately alternating between different stress modalities, can exploit trade-offs between resistance mechanisms and delay the emergence of multi-resistant clones. This distinction between the tumor cell game, played among cancer, stromal, and immune cells with stress-adaptive phenotypes as strategies, and the clinician–tumor metagame, in which the clinician chooses treatment strategies to shift the payoff matrix of the underlying cell game, is conceptually important. The metagame is fundamentally asymmetric in information, as the clinician cannot directly observe the tumor's payoff matrix or strategy frequencies but must infer them from noisy, delayed surrogates such as tumor burden, ctDNA, and imaging. Adaptive therapy, in this light, is not a fixed protocol but an ongoing inference problem in which the clinician iteratively updates beliefs about the tumor's game state and adjusts dosing accordingly.

A complementary therapeutic strategy, proposed by Archetti and Pienta[13], targets the cooperative infrastructure of tumors rather than individual cell phenotypes. Because many hallmarks of cancer depend on the production of diffusible public goods (growth factors, angiogenic signals, immune-suppressive cytokines), therapies that disrupt cooperation can trigger a "tragedy of the commons" at the cellular level. One approach involves engineering autologous defector cells that abolish production of cooperative factors while retaining the ability to benefit from neighbors' contributions. When reintroduced, these defectors free-ride on cooperative neighbors, spreading through clonal selection and progressively eroding the cooperative foundation of collective stress adaptation. Importantly, such cooperation-targeting therapies may be inherently evolution-proof, because a population of pure defectors cannot be re-invaded by cooperators, because the benefit of cooperation is shared while its cost is borne only by producers[13]. Once cooperation collapses, the tumor loses access to essential collective functions without a viable evolutionary path to restore them.

Beyond altering payoff magnitudes, certain therapies fundamentally change the type of game being played, a distinction with profound strategic implications. Immunotherapy, for instance, does not merely increase the cost of immune-visible phenotypes; by unleashing T cell-mediated killing, it transforms a tumor-immune interaction that may have settled into a stable coexistence equilibrium into one where immune evasion becomes the dominant selective pressure, effectively converting a Snowdrift-like game into a Prisoner's Dilemma in which non-evaders are eliminated. This "therapy-as-environment" perspective[24] implies that predicting therapeutic outcomes requires understanding not just how payoffs change quantitatively but whether the qualitative game structure, and therefore the set of possible equilibria, is preserved or destroyed. A distinct category of intervention targets the physical infrastructure of the game itself. In nascent oesophageal tumours, in vivo inhibition of fibronectin fibrillogenesis (FUD peptide) or EGFR signalling (gefitinib) significantly reduced precancerous niche formation[46], collapsing the stromal scaffold on which the tumour–fibroblast cooperative game depends rather than modifying payoffs within it.

A further implication is that therapeutic stress does not only reshape the game within the primary tumor but can inadvertently create new games in distant compartments. Surgery triggers systemic wound-healing responses and inflammatory cascades that can awaken dormant disseminated tumor cells, while even low-dose radiotherapy remodels the pre-metastatic niche by activating neutrophils and modifying lung epithelial cells[56]. These therapy-induced metastatic effects represent unintended

perturbations of games the clinician was not targeting, underscoring the need for game-theoretic models that account for multi-compartment dynamics rather than treating the primary tumor in isolation.

The game-theoretic perspective also explains why combination therapies often fail despite mechanistic rationale. If two drugs target strategies that are mutually exclusive due to trade-offs, simultaneous administration may be less effective than sequential dosing, as cells cannot simultaneously resist both pressures. Conversely, if drugs target strategies that can coexist, simultaneous treatment is required to prevent either strategy from escaping. Predicting which scenario applies requires understanding the structure of phenotypic trade-offs, not merely the molecular targets of each drug.

When Adaptation Becomes Vulnerability

Adaptation improves short-term fitness but imposes constraints. Evolutionary theory predicts that specialization narrows future options, creating trade-offs that can be exploited when environments change. In cancer, adaptations that buffer one form of stress often increase sensitivity to others.

These constraints are not immediately apparent in stable environments but become evident when stresses shift. A canonical example comes from hypoxia-adapted cells in solid tumors. Prolonged hypoxia selects for cells with enhanced glycolytic capacity, suppressed mitochondrial oxidative phosphorylation, and elevated reductive carboxylation for lipid synthesis. These adaptations are protective in low-oxygen conditions but render cells vulnerable to reoxygenation injury. Upon restoration of oxygen, dysregulated mitochondrial metabolism generates excessive ROS, and cells with suppressed antioxidant defenses experience oxidative damage. This phenomenon has been exploited therapeutically. Reoxygenation, in which tumors are transiently exposed to elevated oxygen tension, selectively kills hypoxia-adapted cells while sparing normoxic populations[77]. Similarly, cells that adapt to nutrient limitation through autophagy upregulation become dependent on lysosomal function and are sensitized to lysosomal inhibitors such as chloroquine. In KRAS-mutant pancreatic cancer, where basal autophagy is elevated due to metabolic stress[78], chloroquine shows selective toxicity compared to KRAS-wild-type cells, illustrating how adaptation to one stress creates collateral sensitivity to mechanistically related perturbations. Another well-characterized example involves immune-adapted tumors. Melanomas that escape immune surveillance through constitutive interferon signaling and antigen presentation machinery downregulation become resistant to checkpoint blockade but are often sensitized to MAPK pathway inhibitors, as IFN signaling and MAPK signaling exhibit antagonistic interactions[79]. Sequential treatment strategies exploiting this trade-off have shown promise in preclinical models.

These trade-offs can be formalized through the concept of anti-correlated payoff vectors. If adaptation to stress A shifts a cell's phenotype such that its fitness under stress B decreases proportionally, then the payoff vectors for strategies optimized to A and B are anti-correlated. Sequential therapies and fluctuating treatment schedules exploit this structure by preventing any single strategy from achieving high payoffs across consecutive rounds of the game. The degree of anti-correlation between payoff vectors determines the therapeutic leverage available, with perfectly anti-correlated payoffs (where resistance to drug A implies maximal sensitivity to drug B) create an evolutionary trap from which no mixed strategy can escape, while weakly correlated payoffs permit generalist strategies that partially resist both pressures.

We argue that understanding adaptation as a constrained optimization problem should shift therapeutic thinking from targeting pathways to targeting the limits of adaptability itself. Rather than asking "what pathway is activated," we should ask "what constraints does this adaptation impose, and how can we exploit the anti-correlations it creates?" This logic extends to individual genes themselves. Many canonical cancer genes, TGF-β, p53, Notch, autophagy regulators, function as "double agents" whose payoff contribution reverses sign depending on tumor stage, microenvironmental context, or cellular strategy composition[38]. TGF-β, for instance, suppresses proliferation in early-stage tumors but promotes invasion and immune evasion in advanced disease; the same molecular player switches from imposing costs on defectors to subsidizing them as the game evolves. Recognizing these context-dependent sign reversals is essential for avoiding therapeutic strategies that inadvertently convert a tumor suppressor into an oncogenic signal by altering the game context in which it operates.

A critical assumption underlying many therapeutic strategies, including adaptive therapy, is that resistance mechanisms carry fitness costs in the absence of treatment. However, recent empirical work challenges this assumption. Kaznatcheev and colleagues[24] demonstrated that alectinib-resistant non-small cell lung cancer cells did not exhibit classical costs of resistance; in fact, resistant cells outperformed sensitive cells even in drug-free co-culture conditions. This "negative cost" of resistance fundamentally alters predictions about competitive dynamics and adaptive dosing strategies. The observation underscores the importance of empirically measuring interaction structures rather than relying on intuition, and suggests that some cancers may require fundamentally different evolutionary management strategies than those based on competitive suppression of resistant cells by sensitive cells.

Measuring Stress Games Experimentally

Studying cancer as an evolutionary game requires experimental approaches that capture interactions, competition, and temporal dynamics (**Box 3**). Static endpoint measurements obscure the relative fitness of strategies and fail to reveal how equilibria form or dissolve.

Single-cell approaches for strategy identification

A central challenge in applying game theory to cancer is bridging the gap between the abstract concept of a "strategy" and the measurable biology of individual cells. Single-cell omics technologies have advanced this effort by revealing that tumor cell populations organize into discrete, recurrent transcriptional cell states, clusters of cells occupying distinct regions of transcriptional space, each characterized by coherent gene expression programs associated with specific stress-adaptive functions such as proliferation, invasion, drug tolerance, or immune evasion[25,26]. These cell states are not genetically determined in most cases; cells carrying identical mutations can occupy different states, and individual cells transition between states over time in response to microenvironmental cues. However, cell states are not strategies. A cell state is a transcriptional snapshot: it tells us what a cell is expressing at a given moment. A strategy, by contrast, is a conditional behavioral policy that maps the cell's game state (the local configuration of stresses, resources, and neighbor strategies) to a behavioral output with fitness consequences that depend on what surrounding cells are doing. Two cells in the same transcriptional state, for example, both expressing a hypoxic gene program, may be executing different strategies if their neighborhoods differ: one surrounded by glycolytic neighbors (glucose-depleted game state) may adopt quiescence, while another surrounded by quiescent neighbors (glucose-available game state) may adopt glycolysis. The game-theoretic interpretation therefore requires two steps: first, single-cell profiling identifies recurrent cell states as candidate strategy signatures; second, spatial or contextual information about the game state, neighbor composition, microenvironmental gradients, is needed to infer which strategy each cell state implements in a given context. The relative abundance of cell states in a tumor reflects a proxy for the population's strategy distribution, and transitions between states correspond to strategy switching, but the mapping from state to strategy is context-dependent rather than one-to-one (**Fig. 4A**).

Longitudinal single-cell tracking is essential for directly observing these strategy dynamics under controlled stress conditions. Time-lapse microscopy coupled with fluorescent reporters enables measurement of individual cell fate decisions, proliferation rates, and phenotypic transitions in real time. When combined with environmental perturbations such as drug addition, nutrient depletion, or reoxygenation, these approaches can quantify how payoffs shift and which strategies succeed under different conditions.

Single-cell RNA sequencing, particularly when performed at multiple time points during stress exposure, provides high-resolution snapshots of cell states and allows inference of trajectory dynamics. RNA velocity analysis can predict future cell states based on the balance of unspliced and spliced transcripts, enabling reconstruction of stress-response trajectories without direct lineage tracing. Recent applications in melanoma have identified at least four recurrent cell states, proliferative, invasive, neural crest-like, and a transient drug-tolerant phenotype[25,64], with evidence of stochastic transitions between states at baseline and directional transitions under MAPK inhibition. Crucially, the drug-tolerant state is not a pre-existing genetic clone but a reversible transcriptional program that cells enter under therapeutic stress and exit upon drug withdrawal, consistent with the game-theoretic prediction that strategy switching, rather than clonal selection alone, mediates early resistance. These cell states serve as the empirical entry point for strategy inference: identifying which strategy each state implements requires overlaying game state information, the local stress regime and neighbor composition, onto the transcriptional data. Similar cell state architectures have been identified across cancer types, including a mesenchymal-like drug-tolerant persister state in lung adenocarcinoma and a quiescent stem-like state in glioblastoma, suggesting that the repertoire of stress-adaptive cell states from which strategies are drawn may be convergent across malignancies.

Integrating single-cell ATAC-seq with RNA-seq provides insight into the regulatory logic governing strategy switching by revealing which enhancers and transcription factors control access to alternative cell states. In the context of therapy resistance, such multi-omic approaches have identified pioneer transcription factors that reprogram chromatin accessibility and enable entry into drug-tolerant states. For example, joint profiling of chromatin accessibility and gene expression during the emergence of drug tolerance in non-small cell lung cancer has revealed that the transition from a proliferative to a persister cell state involves coordinated opening of AP-1 family transcription factor binding sites and closure of lineage-specifying enhancers, establishing a distinct epigenetic configuration that is poised yet reversible. Perturbation-based single-cell methods such as Perturb-seq, which combine pooled CRISPR perturbations with single-cell RNA sequencing readout, enable systematic mapping of how individual gene knockouts alter cell state distributions and transition probabilities under stress, effectively dissecting the genetic control of strategy adoption at single-cell resolution.

Lineage tracing and clonal competition assays

Lineage tracing using DNA barcoding or genome editing-based recorders allows direct measurement of clonal fitness over time. By tagging cells with unique heritable barcodes and tracking barcode frequencies through serial sequencing, one can infer relative growth rates, survival under stress, and competitive interactions. Pooled CRISPR screens represent a specialized form of this approach, in which genetic perturbations are coupled to barcodes, enabling systematic exploration of how specific genes influence strategy adoption and fitness under different stress conditions.

Crucially, CRISPR screens under fluctuating or sequential stress conditions can identify genes controlling phenotypic plasticity rather than constitutive stress resistance. For instance, performing dropout screens under cycling hypoxia reveals genes required for adaptation to fluctuating oxygen, which differ from those required for chronic hypoxia tolerance. Similarly, screens performed during drug treatment followed by drug withdrawal can identify genes controlling reversible drug tolerance, distinguishing them from genes mediating irreversible genetic resistance.

Competitive mixing experiments, in which defined subpopulations are co-cultured at varying ratios and their relative abundances tracked over time, directly test game-theoretic predictions about frequency-dependent fitness. These assays can be performed with genetically or epigenetically distinct subclones, or with cells sorted based on phenotypic markers. By varying the initial frequency of each strategy and the environmental conditions, one can empirically construct payoff matrices and test whether observed dynamics match game-theoretic equilibria (**Fig. 4B**).

Spatial and microenvironmental profiling

Spatial transcriptomics and multiplexed imaging technologies provide in situ characterization of cell state distributions within tumors, revealing how spatial structure and microenvironmental gradients influence the coexistence of distinct adaptive strategies. Technologies such as MERFISH, Slide-seq, and Visium enable simultaneous measurement of hundreds to thousands of transcripts while preserving spatial coordinates, allowing researchers to ask whether specific cell states are enriched in particular microenvironmental niches. In colorectal cancer, spatial transcriptomics has revealed that proliferative cell states cluster near functional vasculature while invasive, mesenchymal-like states localize to the hypoxic tumor edge, and drug-tolerant states occupy intermediate zones of metabolic stress. By correlating cell state identities with local oxygen tension, nutrient availability, immune infiltration, and stromal architecture, these approaches enable direct mapping of the relationship between the local stress landscape and the strategies cells adopt within it (**Fig. 4C**).

Combining spatial profiling with perturbation, such as targeted ablation of specific cell types or pharmacological modulation of microenvironmental stresses, allows causal testing of how niche interactions shape the stability of cell state distributions. For example, depleting cancer-associated fibroblasts in pancreatic cancer disrupts metabolic symbiosis and shifts the balance between glycolytic and OXPHOS tumor cell states, consistent with game-theoretic predictions about the role of stromal cells in maintaining heterogeneity. Similarly, spatial profiling before and after immunotherapy has revealed that responding tumors undergo wholesale reorganization of cell state composition, with immune-evasive states collapsing and proliferative states expanding, while resistant tumors maintain stable cell state distributions that reflect pre-existing equilibria robust to immune perturbation.

Stress-response reporter systems for real-time game readout

The experimental approaches described above, single-cell profiling, lineage tracing, spatial transcriptomics, provide high-resolution but largely retrospective snapshots of strategy composition. A critical missing layer is the ability to monitor stress-adaptive strategy adoption in real time and at scale. Plasmid-based transcriptional reporter systems offer a direct route to this capability by converting the activation of specific stress-response programs into fluorescent or luminescent readouts that can be tracked continuously in living cells.

The design principle is straightforward: a minimal stress-responsive promoter element drives expression of a fluorescent protein, such that reporter intensity serves as a quantitative proxy for pathway activation. Reporters for the major stress axes discussed in this Review are already available or readily constructable. Hypoxia-responsive elements (HREs) driving GFP report HIF pathway activation and identify cells adopting hypoxic survival strategies[80]. Antioxidant response elements (AREs) driving mCherry or similar fluorophores read out NRF2-mediated oxidative stress adaptation. Unfolded protein response elements (UPREs) coupled to fluorescent reporters monitor ER stress and proteotoxic adaptation, while heat shock elements (HSEs) report HSF1-dependent proteostasis programs. DNA damage reporters based on p21 or GADD45A promoter elements can identify cells engaging checkpoint activation or damage tolerance strategies[81,82]. Each reporter maps directly onto a strategy axis in the game-theoretic framework: a cell activating the HRE reporter is adopting a hypoxic tolerance strategy, while one activating the ARE reporter is investing in oxidative stress defense, and so on.

The power of this approach multiplies when reporters for orthogonal stress axes are combined in the same cell line through spectral multiplexing. A dual-reporter system (e.g., HRE::GFP + ARE::mCherry) enables simultaneous readout of two stress-response programs, partitioning cells into four strategy quadrants, double-negative, HIF-only, NRF2-only, and double-positive, whose relative frequencies can be tracked by flow cytometry or live-cell imaging over time. Extending to

three or four reporters using spectrally distinct fluorophores (BFP, GFP, YFP, mCherry) or orthogonal readout modalities (fluorescence plus luminescence) would enable real-time monitoring of strategy distributions across the major stress axes simultaneously. Such multiplexed reporter panels would, for the first time, allow direct visualization of how a tumor cell population redistributes across the strategy space in response to a therapeutic perturbation, the empirical equivalent of watching the game switch in real time.

Critically, these reporters can be integrated with the competition assays described above. In a mixing experiment with reporter-labeled cells, one can measure not only which subpopulation expands or contracts (the payoff) but also which stress programs are active in each subpopulation and how activation patterns shift with changing frequency composition. This enables a richer payoff matrix that captures not just growth rate as a function of frequency but growth rate as a function of both frequency and strategy state, linking fitness directly to the stress-adaptive programs that define each strategy. Reporter-equipped cells can also be used in CRISPR screens to identify genes whose knockout shifts the balance between stress strategies, or in spatial co-culture assays to map how local microenvironmental context determines which reporters activate in which cells.

A practical consideration is that plasmid-based reporters offer rapid deployment, transfection or transduction of reporter constructs into existing cell line models requires days rather than months, and are compatible with standard flow cytometry, live-cell imaging, and plate-reader workflows already established in most cancer biology laboratories. This low barrier to adoption could accelerate the transition from retrospective to real-time game measurement across the field.

Computational inference of payoff matrices

A major challenge in applying game theory to cancer is the empirical estimation of payoff matrices. Standard approaches in evolutionary biology involve controlled competition experiments, but these are labor-intensive and difficult to scale. The pioneering work of Kaznatcheev and colleagues[24] established a game assay for cancer by combining co-culture experiments at varying initial frequencies with time-lapse microscopy to measure growth rates as proxies for fitness. By fitting these growth rate versus frequency data to fitness functions, they could infer the payoff matrix and classify the game being played. Importantly, they distinguished between reductive games (interactions at the level of individual cell pairs) and effective games (population-level dynamics that emerge from spatial structure and ecological context), recognizing that spatial heterogeneity can transform one game type into another.

Recent computational methods leverage time-series single-cell data to infer fitness landscapes and interaction terms through maximum likelihood or Bayesian inference frameworks, extending the game assay approach to higher-dimensional phenotypic spaces. One approach involves fitting Lotka-Volterra competition models to longitudinal subpopulation frequencies, parameterizing intrinsic growth rates and interaction coefficients that can be interpreted as game-theoretic payoffs. However, honest assessment reveals fundamental limitations. Payoff matrices estimated from in vitro co-culture may not transfer to in vivo contexts where spatial structure, immune interactions, and systemic factors reshape effective fitness; multiple distinct payoff matrices can generate nearly identical short-term dynamics, creating severe identifiability problems from sparse clinical measurements; and the discrete-strategy assumption central to payoff matrix estimation may inadequately capture phenotypic continuity in some tumor contexts. These limitations do not invalidate the approach but define its current boundaries and motivate complementary methods.

Machine learning approaches, including neural differential equations and variational autoencoders, offer data-driven routes to constructing payoff matrices by learning latent cell state representations and inferring transition probabilities from time-series single-cell data. Variational inference frameworks can decompose observed cell state dynamics into contributions from differential proliferation (selection on existing strategies), directed state transitions (strategy switching), and stochastic noise, enabling researchers to distinguish whether compositional shifts under stress reflect clonal selection or active phenotypic switching.

Integrating experimental and computational approaches

The most powerful approach combines experimental perturbation with computational modeling in an iterative cycle. Initial single-cell profiling identifies recurrent cell states through clustering or trajectory analysis, providing candidate signatures from which strategies can be inferred once game state information is overlaid. Competition assays and stress perturbations test hypotheses about fitness and plasticity. Computational models integrate these data to infer payoffs and predict dynamics under untested conditions. Predictions are then validated experimentally, and discrepancies refine the model. This cycle mirrors experimental evolution studies in microbiology, which have successfully dissected the genetic and ecological basis of adaptation. In cancer, such integrated approaches remain rare but are increasingly feasible. The combination of CRISPR screening, single-cell multi-omics, spatial profiling, and computational inference provides the *images and words*—spatial maps and transcriptional profiles, needed to study cancer as an evolving game in real time.

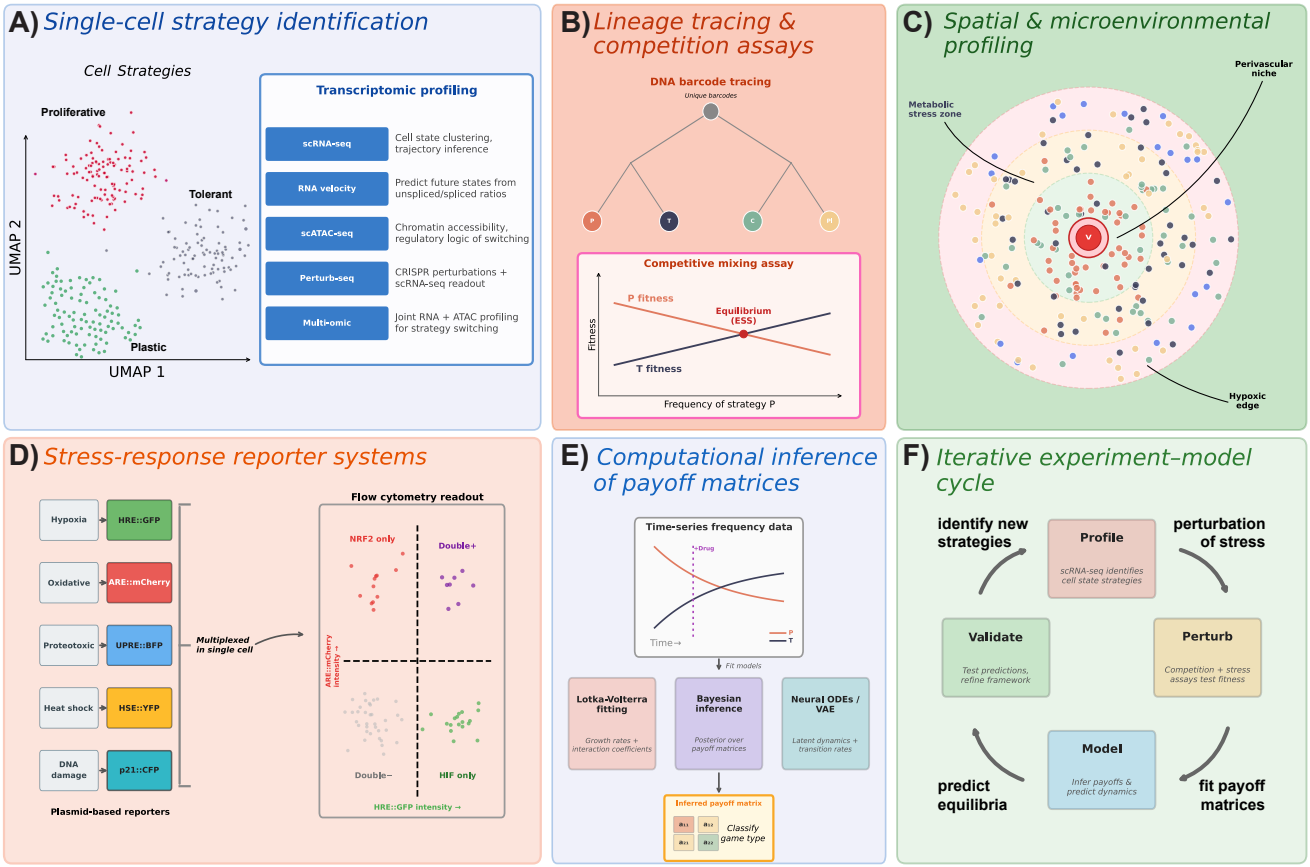

Figure 4. Experimental approaches for measuring stress games in cancer. (A) From cell states to strategy inference. Tumor cells cluster into discrete transcriptional cell states on UMAP projections; these states provide observable signatures from which game-theoretic strategies are inferred by integrating spatial and microenvironmental context (the game state). Key technologies include scRNA-seq, RNA velocity, single-cell ATAC-seq, Perturb-seq, spatial transcriptomics, and joint multi-omic profiling. **(B)** Lineage tracing and competition assays. DNA barcode tracing infers clonal fitness through serial sequencing. Competitive mixing assays reveal frequency-dependent fitness, with the intersection of strategy fitness curves marking the ESS equilibrium. **(C)** Spatial and microenvironmental profiling. A tumor cross-section shows concentric organization relative to vasculature (V), with proliferative cells (orange) in perivascular niches, tolerant and cooperative cells (dark grey, green) in metabolic stress zones, and plastic and dormant cells (yellow, blue) at the hypoxic edge. **(D)** Stress-response reporter systems. Plasmid-based transcriptional reporters (HRE::GFP, ARE::mCherry, UPRE::BFP, HSE::YFP, p21::CFP) convert stress-pathway activation into fluorescent readouts. Spectral multiplexing in a single cell line enables simultaneous monitoring of multiple stress axes, with flow cytometry partitioning cells into strategy quadrants (e.g., HIF-only, NRF2-only, double-positive, double-negative) whose frequencies can be tracked over time during therapeutic perturbation. **(E)** Computational inference of payoff matrices. Time-series frequency data from competition assays or longitudinal single-cell profiling are fitted using Lotka-Volterra models, Bayesian inference, or neural ODE/VAE frameworks to estimate interaction coefficients and infer the payoff matrix governing the evolutionary game. **(F)** Iterative experiment–model cycle. Single-cell profiling identifies candidate strategies, perturbation assays test fitness and plasticity, computational models infer payoffs and predict dynamics under untested conditions, and experimental validation refines the framework. This cycle integrates CRISPR screens, multi-omics, spatial profiling, reporter readouts, and machine learning inference into a unified workflow for studying cancer as an evolving game.

From Evolutionary Logic to Therapeutic Strategy

We take the position that if tumor heterogeneity, plasticity, and resistance are predictable outcomes of selection under constraint, then therapy should be designed to manipulate the game rather than simply eliminate players. Four principles follow directly from the framework developed above, each targeting a different feature of the evolutionary dynamics that sustain tumors.

First, stress adaptation imposes costs that create exploitable trade-offs. Cells that invest in DNA damage tolerance sacrifice proliferative speed; cells that upregulate immune evasion programs divert resources from metabolic efficiency. These trade-offs generate collateral sensitivities, such that a population driven into one adaptive strategy becomes newly vulnerable to treatments targeting the capacities it has sacrificed. Enriquez-Navas and colleagues demonstrated this principle by showing that evolutionary steering, deliberately selecting for one resistant phenotype, can restore sensitivity to a previously failed drug[83]. Translating this insight requires sequencing therapies not by toxicity profiles but by the adaptive trade-offs they impose, so that each treatment round opens vulnerabilities exploited by the next. To illustrate with a concrete logic chain. In the glycolytic–OXPHOS Snowdrift game described above, anti-angiogenic therapy intensifies hypoxia, increasing the payoff to glycolysis and shifting the equilibrium toward glycolytic dominance. At this new equilibrium, the population becomes collectively vulnerable to agents such as dichloroacetate (DCA) that force pyruvate into mitochondrial oxidation, a metabolic demand the now-dominant glycolytic cells are poorly equipped to sustain. The prediction is directly testable. Anti-angiogenic pre-treatment followed by DCA should yield greater cytotoxicity than either agent alone, and the magnitude of benefit should scale with the degree of glycolytic enrichment achieved during the pre-treatment phase, measurable by lactate-to-pyruvate ratio or ^{18}F-FDG-PET uptake.

Second, phenotypic plasticity enables cells to escape therapeutic pressure by transitioning between cell states, and thereby switching strategies, rather than acquiring new mutations. Inhibiting the molecular machinery of state switching, chromatin remodelers, pioneer transcription factors, stress-responsive epigenetic modifiers, can lock tumor cells into a single cell state, restricting the strategies accessible to them and rendering them vulnerable to strategy-specific attack[84,85]. This reframes drug resistance as a problem of restricting accessible phenotypic and strategic space rather than targeting individual resistance mechanisms one at a time.

Third, the adaptive therapy paradigm proposes maintaining a population of drug-sensitive cells that competitively suppress resistant clones through frequency-dependent fitness interactions, accepting containment rather than eradication as the therapeutic objective[86]. The pilot trial of adaptive abiraterone dosing in metastatic castration-resistant prostate cancer reported extended time to progression by modulating dose to sustain competition between sensitive and resistant populations[76], and a follow-up study confirmed that evolution-informed mathematical models could further improve outcomes[87]. Whether the observed benefit reflects genuine competitive suppression or altered drug exposure dynamics remains debated, and a systematic survey of open questions highlights the absence of prospective randomized trials and validated real-time biomarkers of clonal composition[88]. Nonetheless, the principle that maintaining ecological balance within tumors may outperform maximum tolerated dose regimens has motivated clinical investigation across multiple cancer types[89].

Fourth, the spatial organization of tumors sustains strategy coexistence by creating microenvironmental niches that locally favor different phenotypes. Treatments that disrupt this architecture, anti-angiogenic agents that redistribute oxygen gradients, stromal-targeted therapies that alter paracrine signaling, or immunotherapies that reorganize the immune landscape, can collapse the spatial heterogeneity that maintains game-theoretic equilibria[90]. A tumor driven from a spatially structured Snowdrift game toward a well-mixed Prisoner's Dilemma loses its capacity for stable coexistence and becomes more susceptible to conventional cytotoxic approaches.

Translating these principles to the clinic faces three concrete barriers. Measuring payoff matrices in individual patients remains impractical, as ex vivo competition assays on patient-derived organoids can estimate payoff coefficients, but these measurements omit immune selection, vascular dynamics, and systemic factors that reshape effective fitness in vivo. Real-time readouts of strategy composition are lacking, as circulating tumor DNA and imaging provide partial views of clonal architecture but neither maps onto game-theoretic state variables with the temporal resolution required for adaptive dosing. Multiple distinct payoff structures can produce nearly identical short-term dynamics, creating a fundamental identifiability problem for model-based treatment optimization[91,92]. These barriers are substantial but increasingly tractable. Patient-derived organoid avatars from serial biopsies enable personalized fitness estimation without repeated invasive sampling. Bayesian frameworks integrating longitudinal liquid biopsy data with mechanistic models can progressively refine payoff estimates as measurements accumulate. Reinforcement learning and optimal control methods offer dosing strategies that are robust to parameter uncertainty[91]. The challenge ahead is not whether evolutionary game theory can inform cancer therapy, but whether the measurement and modeling infrastructure required for bedside implementation can be built fast enough to keep pace with the theory.

Conclusion

The central difficulty of cancer is not that it resists treatment but that it adapts. Every selective pressure reshapes the fitness landscape and favors cells whose strategies best exploit the new equilibrium. The framework developed here repositions this adaptability from a clinical frustration to a predictable, quantifiable, and potentially targetable feature of tumor biology. Stress adaptation, we propose, is not one pathway among many but the evolutionary logic that organizes heterogeneity, drives plasticity, and gives rise to resistance.

Evolutionary game theory offers a formal language well suited to this challenge. Cooperation, competition, and coexistence among cancer cell phenotypes are not merely metaphors: they are measurable outcomes of frequency-dependent selection operating on identifiable strategies. We propose that therapies do not merely select for resistance but can switch the type of game being played. The empirical demonstration of this principle, from Prisoner's Dilemma to Snowdrift, from Leader to Deadlock, reframes treatment from a problem of elimination to one of strategic redirection. The six predictions advanced in this review are testable with technologies that already exist: single-cell multi-omics to identify strategies, lineage tracing to measure fitness, spatial transcriptomics to map the arenas where games are played, and computational inference to estimate the payoff matrices that govern outcomes.

Important open questions accompany this vision. Payoff matrices measured in organoids may not transfer to the immunological and spatial complexity of tumors in vivo. The discrete-strategy approximation, while motivated by recurrent cell states in single-cell data, faces the additional challenge that inferring strategies from cell states requires game state information, neighbor composition and microenvironmental context, that remains difficult to obtain at scale; moreover, the approximation may obscure continuous phenotypic gradients in some contexts. Clinical validation extends little beyond a single adaptive therapy trial in prostate cancer[76] and a handful of exploratory studies; prospective trials designed to test whether competitive suppression, rather than pharmacokinetic artifact, mediates the benefit of adaptive dosing are urgently needed. These are not reasons for skepticism. They are a research agenda.

There is reason for optimism, however. Single-cell genomics, spatial profiling, patient-derived organoid systems, and machine learning methods for inferring dynamics from sparse measurements are each, independently, reaching the resolution required to parameterize evolutionary games in real tumors. Their integration will determine whether the next decade of oncology continues to treat cancer as a collection of molecular targets or engages with it as an evolving ecology of competing strategies, shaped by stress, sustained by trade-offs, and governed by rules we are only now learning to read. As these tools mature, the game-theoretic perspective offers a unifying language for synthesizing observations across scales and modalities, transforming cancer biology from a catalog of pathways into a predictive science of adaptation under constraint. The alternative, continuing to target molecular pathways while hoping the next combination will outpace adaptation—is *running on faith* in a game that rewards only those who learn its rules.

Data Availability

No datasets were generated or analyzed during the current study.

Author Contributions

S.C.O. conceived and wrote the manuscript.

Acknowledgments

The author thanks Mert Gumren, a behavioural economist, for conversations on game theory and strategic decision-making years before this manuscript took shape, which fundamentally influenced the author's thinking in game-theoretic terms and inspired the interdisciplinary perspective of this work. The author also thanks Itai Yanai for encouraging the author to continue working in this area.

Competing Interests

The author declares no competing interests.

References

1. Lior, C. *et al.* Mapping the tumor stress network reveals dynamic shifts in the stromal oxidative stress response. *Cell Reports* **43** (2024).

2. Dominguez-Sola, D. *et al.* Non-transcriptional control of DNA replication by c-Myc. *Nature* **448**, 445–451, DOI: 10.1038/nature05953 (2007).

3. DeNicola, G. M. *et al.* Oncogene-induced Nrf2 transcription promotes ROS detoxification and tumorigenesis. *Nature* **475**, 106–109, DOI: 10.1038/nature10189 (2011).

4. Stylianopoulos, T. *et al.* Causes, consequences, and remedies for growth-induced solid stress in murine and human tumors. *Proc. Natl. Acad. Sci.* **109**, 15101–15108, DOI: 10.1073/pnas.1213353109 (2012).

5. Zaretsky, J. M. *et al.* Mutations associated with acquired resistance to PD-1 blockade in melanoma. *New Engl. J. Medicine* **375**, 819–829, DOI: 10.1056/NEJMoa1604958 (2016).

6. Wang, H. *et al.* Cancer neuroscience in head and neck: interactions, modulation, and therapeutic strategies. *Mol. Cancer* **24**, 101, DOI: 10.1186/s12943-025-02299-6 (2025).

7. Geller, L. T. *et al.* Potential role of intratumor bacteria in mediating tumor resistance to the chemotherapeutic drug gemcitabine. *Science* **357**, 1156–1160, DOI: 10.1126/science.aah5043 (2017).

8. Levy, M., Blacher, E. & Elinav, E. Microbiome, metabolites and host immunity. *Curr. Opin. Microbiol.* **35**, 8–15, DOI: 10.1016/j.mib.2016.10.003 (2016).

9. Gatenby, R. A. & Brown, J. S. The evolution and ecology of resistance in cancer therapy. *Cold Spring Harb. Perspectives Medicine* **8**, a033415, DOI: 10.1101/cshperspect.a033415 (2018).

10. Kareva, I., Luddy, K. A., O'Farrelly, C., Gatenby, R. A. & Brown, J. S. Predator-prey in tumor-immune interactions: A wrong model or just an incomplete one? *Front. Immunol.* **12**, 668221, DOI: 10.3389/fimmu.2021.668221 (2021).

11. Axelrod, R., Axelrod, D. E. & Pienta, K. J. Evolution of cooperation among tumor cells. *Proc. Natl. Acad. Sci.* **103**, 13474–13479, DOI: 10.1073/pnas.0606053103 (2006).

12. Archetti, M., Ferraro, D. A. & Christofori, G. Heterogeneity for igf-ii production maintained by public goods dynamics in neuroendocrine pancreatic cancer. *Proc. Natl. Acad. Sci.* **112**, 1833–1838, DOI: 10.1073/pnas.1414653112 (2015).

13. Archetti, M. & Pienta, K. J. Cooperation among cancer cells: applying game theory to cancer. *Nat. Rev. Cancer* **19**, 110–117, DOI: 10.1038/s41568-018-0083-7 (2019).

14. Wright, S. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In *Proceedings of the Sixth International Congress of Genetics*, vol. 1, 356–366 (1932).

15. Kauffman, S. A. *The Origins of Order: Self-Organization and Selection in Evolution* (Oxford University Press, New York, 1993).

16. Dieckmann, U. & Law, R. The dynamical theory of coevolution: a derivation from stochastic ecological processes. *J. Math. Biol.* **34**, 579–612, DOI: 10.1007/BF02409751 (1996).

17. Metz, J. A. J., Geritz, S. A. H., Meszéna, G., Jacobs, F. J. A. & van Heerwaarden, J. S. Adaptive dynamics: A geometrical study of the consequences of nearly faithful reproduction. In van Strien, S. J. & Verduyn Lunel, S. M. (eds.) *Stochastic and Spatial Structures of Dynamical Systems*, 183–231 (North-Holland, Amsterdam, 1996).

18. Bozic, I. *et al.* Evolutionary dynamics of cancer in response to targeted combination therapy. *eLife* **2**, e00747, DOI: 10.7554/eLife.00747 (2013).

19. Nowak, M. A. *Evolutionary Dynamics: Exploring the Equations of Life* (Harvard University Press, Cambridge, MA, 2006).

20. França, G. S. *et al.* Cellular adaptation to cancer therapy along a resistance continuum. *Nature* **631**, 876–883 (2024).

21. Sharma, S. V. *et al.* A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. *Cell* **141**, 69–80, DOI: 10.1016/j.cell.2010.02.027 (2010).

22. Stanková, K., Brown, J. S., Dalton, W. S. & Gatenby, R. A. Cancer heterogeneity and multilayer spatial evolutionary games. *Biol. Direct* **11**, 52, DOI: 10.1186/s13062-016-0156-z (2016).

23. Fassl, A., Geng, Y. & Sicinski, P. CDK4 and CDK6 kinases: From basic science to cancer therapy. *Science* **375**, eabc1495, DOI: 10.1126/science.abc1495 (2022).

24. Kaznatcheev, A., Peacock, J., Basanta, D., Marusyk, A. & Scott, J. G. Fibroblasts and alectinib switch the evolutionary games played by non-small cell lung cancer. *Nat. Ecol. & Evol.* **3**, 450–456, DOI: 10.1038/s41559-018-0768-z (2019).

25. Tirosh, I. *et al.* Dissecting the multicellular ecosystem of metastatic melanoma by single-cell RNA-seq. *Science* **352**, 189–196, DOI: 10.1126/science.aad0501 (2016).

26. Gavish, A. *et al.* Hallmarks of transcriptional intratumour heterogeneity across a thousand tumours. *Nature* **618**, 598–606, DOI: 10.1038/s41586-023-06130-4 (2023).

27. Lord, C. J. & Ashworth, A. Mechanisms of resistance to therapies targeting BRCA-mutant cancers. *Nat. Medicine* **19**, 1381–1388, DOI: 10.1038/nm.3369 (2013).

28. Rondinelli, B. *et al.* EZH2 promotes degradation of stalled replication forks by recruiting MUS81 through histone H3 trimethylation. *Nat. Cell Biol.* **19**, 1371–1378, DOI: 10.1038/ncb3626 (2017).

29. Commisso, C. *et al.* Macropinocytosis of protein is an amino acid supply route in Ras-transformed cells. *Nature* **497**, 633–637, DOI: 10.1038/nature12138 (2013).

30. Schewe, D. M. & Aguirre-Ghiso, J. A. ATF6α-Rheb-mTOR signaling promotes survival of dormant tumor cells in vivo. *Proc. Natl. Acad. Sci. USA* **105**, 10519–10524, DOI: 10.1073/pnas.0800939105 (2008).

31. Ratnikov, B. *et al.* Glutamate and asparagine cataplerosis underlie glutamine addiction in melanoma. *Oncotarget* **6**, 7379–7389, DOI: 10.18632/oncotarget.3132 (2015).

32. Lim, J.-H., Luo, C., Vazquez, F. & Puigserver, P. Targeting mitochondrial oxidative metabolism in melanoma causes metabolic compensation through glucose and glutamine utilization. *Cancer Res.* **74**, 3535–3545, DOI: 10.1158/0008-5472.CAN-13-2893 (2014).

33. Kreuzaler, P., Panina, Y., Segal, J. & Yuneva, M. Adapt and conquer: Metabolic flexibility in cancer growth, invasion and evasion. *Mol. Metab.* **33**, 83–101, DOI: 10.1016/j.molmet.2019.08.002 (2020).

34. Duraj, T. *et al.* Beyond the Warburg effect: Oxidative and glycolytic phenotypes coexist within the metabolic heterogeneity of glioblastoma. *Cells* **10**, 202, DOI: 10.3390/cells10020202 (2021).

35. Wölfl, B., te Rietmole, H., Stanková, K., Brown, J. S. & Gatenby, R. A. The contribution of evolutionary game theory to understanding and treating cancer. *Dyn. Games Appl.* **12**, 313–342, DOI: 10.1007/s13235-021-00397-w (2022).

36. Archetti, M. Evolutionary game theory of growth factor production: implications for tumour heterogeneity and resistance to therapies. *Br. J. Cancer* **109**, 1056–1062, DOI: 10.1038/bjc.2013.336 (2013).

37. Salimi Sartakhti, J., Manshaei, M. H., Basanta, D. & Sadeghi, M. Evolutionary emergence of angiogenesis in avascular tumors using a spatial public goods game. *PLoS ONE* **12**, e0175063, DOI: 10.1371/journal.pone.0175063 (2017).

38. bin Masroni, M. S. *et al.* Sociobiology meets oncology: unraveling altruistic cooperation in cancer cells and its implications. *Exp. & Mol. Medicine* **57**, 30–40, DOI: 10.1038/s12276-024-01387-9 (2025).

39. Avraham-Davidi, I. *et al.* Spatially defined multicellular functional units in colorectal cancer revealed from single cell and spatial transcriptomics. *eLife* **14**, e104815, DOI: 10.7554/eLife.104815 (2025).

40. Xiao, J. *et al.* Integrating spatial and single-cell transcriptomics reveals tumor heterogeneity and intercellular networks in colorectal cancer. *Cell Death & Dis.* **15**, 326, DOI: 10.1038/s41419-024-06598-6 (2024).

41. Iriondo, O. *et al.* Hypoxic memory mediates prolonged tumor-intrinsic type I interferon suppression to promote breast cancer progression. *Cancer Res.* **84**, 3141–3157, DOI: 10.1158/0008-5472.CAN-23-2028 (2024).

42. Prickaerts, P. *et al.* Hypoxia increases genome-wide bivalent epigenetic marking by specific gain of H3K27me3. *Epigenetics & Chromatin* **9**, 46, DOI: 10.1186/s13072-016-0086-0 (2016).

43. Pianka, E. R. On r- and k-selection. *The Am. Nat.* **104**, 592–597, DOI: 10.1086/282697 (1970).

44. Aktipis, C. A., Boddy, A. M., Gatenby, R. A., Brown, J. S. & Maley, C. C. Life history trade-offs in cancer evolution. *Nat. Rev. Cancer* **13**, 883–892, DOI: 10.1038/nrc3606 (2013).

45. Amend, S. R. & Pienta, K. J. Ecology meets cancer biology: the cancer swamp promotes the lethal cancer phenotype. *Oncotarget* **6**, 9669–9678, DOI: 10.18632/oncotarget.3430 (2015).

46. Skrupskelyte, G. *et al.* Precancerous niche remodelling dictates nascent tumour persistence. *Nature* DOI: 10.1038/s41586-026-10157-8 (2026). Advance online publication.

47. Yu, T. *et al.* *Fusobacterium nucleatum* promotes chemoresistance to colorectal cancer by modulating autophagy. *Cell* **170**, 548–563.e16, DOI: 10.1016/j.cell.2017.07.008 (2017).

48. Ridlon, J. M., Kang, D. J., Hylemon, P. B. & Bajaj, J. S. Bile acids and the gut microbiome. *Curr. Opin. Gastroenterol.* **30**, 332–338, DOI: 10.1097/MOG.0000000000000057 (2014).

49. Schreiber, R. D., Old, L. J. & Smyth, M. J. Cancer immunoediting: integrating immunity's roles in cancer suppression and promotion. *Science* **331**, 1565–1570, DOI: 10.1126/science.1203486 (2011).

50. Dunn, G. P., Old, L. J. & Schreiber, R. D. The immunobiology of cancer immunosurveillance and immunoediting. *Immunity* **21**, 137–148, DOI: 10.1016/j.immuni.2004.07.017 (2004).

51. Aktipis, C. A. *et al.* Cancer across the tree of life: cooperation and cheating in multicellularity. *Philos. Transactions Royal Soc. B: Biol. Sci.* **370**, 20140219, DOI: 10.1098/rstb.2014.0219 (2015).

52. Wherry, E. J. T cell exhaustion. *Nat. Immunol.* **12**, 492–499, DOI: 10.1038/ni.2035 (2011).

53. Murray, P. J. *et al.* Macrophage activation and polarization: nomenclature and experimental guidelines. *Immunity* **41**, 14–20, DOI: 10.1016/j.immuni.2014.06.008 (2014).

54. Chen, D. S. & Mellman, I. Elements of cancer immunity and the cancer–immune set point. *Nature* **541**, 321–330, DOI: 10.1038/nature21349 (2017).

55. McGranahan, N. *et al.* Clonal neoantigens elicit T cell immunoreactivity and sensitivity to immune checkpoint blockade. *Science* **351**, 1463–1469, DOI: 10.1126/science.aaf1490 (2016).

56. Grasset, E. M., Barillé-Nion, S. & Juin, P. P. Stress in the metastatic journey – the role of cell communication and clustering in breast cancer progression and treatment resistance. *Dis. Model. & Mech.* **17**, dmm050542, DOI: 10.1242/dmm.050542 (2024).

57. Coppé, J.-P., Desprez, P.-Y., Krtolica, A. & Campisi, J. The senescence-associated secretory phenotype: the dark side of tumor suppression. *Annu. Rev. Pathol.* **5**, 99–118, DOI: 10.1146/annurev-pathol-121808-102144 (2010).

58. Coppé, J.-P. *et al.* Senescence-associated secretory phenotypes reveal cell-nonautonomous functions of oncogenic RAS and the p53 tumor suppressor. *PLoS Biol.* **6**, 2853–2868, DOI: 10.1371/journal.pbio.0060301 (2008).

59. Zhang, D. *et al.* Therapy-induced senescent tumor cell-derived extracellular vesicles promote colorectal cancer progression through SERPINE1-mediated NF-κB p65 nuclear translocation. *Mol. Cancer* **23**, 70, DOI: 10.1186/s12943-024-01985-1 (2024).

60. Lopes-Paciencia, S. *et al.* The senescence-associated secretory phenotype and its regulation. *Cytokine* **117**, 15–22, DOI: 10.1016/j.cyto.2019.01.013 (2019).

61. Martinez-Outschoorn, U. *et al.* Oncogenes and inflammation rewire host energy metabolism in the tumor microenvironment: RAS and NFκB target stromal MCT4. *Cell Cycle* **12**, 2580–2597, DOI: 10.4161/cc.25510 (2013).

62. Wilde, L. *et al.* Metabolic coupling and the Reverse Warburg Effect in cancer: Implications for novel biomarker and anticancer agent development. *Semin. Oncol.* **44**, 198–203, DOI: 10.1053/j.seminoncol.2017.10.004 (2017).

63. Guzelsoy, G. *et al.* Cooperative nutrient scavenging is an evolutionary advantage in cancer. *Nature* **640**, 534–542, DOI: 10.1038/s41586-025-08588-w (2025).

64. Tsoi, J. *et al.* Multi-stage differentiation defines melanoma subtypes with differential vulnerability to drug-induced iron-dependent oxidative stress. *Cancer Cell* **33**, 890–904.e5, DOI: 10.1016/j.ccell.2018.03.017 (2018).

65. Beltran, H. *et al.* Divergent clonal evolution of castration-resistant neuroendocrine prostate cancer. *Nat. Medicine* **22**, 298–305, DOI: 10.1038/nm.4045 (2016).

66. Hangauer, M. J. *et al.* Drug-tolerant persister cancer cells are vulnerable to GPX4 inhibition. *Nature* **551**, 247–250, DOI: 10.1038/nature24297 (2017).

67. Botero, C. A., Weissing, F. J., Wright, J. & Rubenstein, D. R. Evolutionary tipping points in the capacity to adapt to environmental change. *Proc. Natl. Acad. Sci.* **112**, 184–189, DOI: 10.1073/pnas.1408589111 (2015).

68. Bakhoum, S. F. *et al.* Chromosomal instability drives metastasis through a cytosolic DNA response. *Nature* **553**, 467–472, DOI: 10.1038/nature25432 (2018).

69. Watkins, T. B. K. *et al.* Pervasive chromosomal instability and karyotype order in tumour evolution. *Nature* **587**, 126–132, DOI: 10.1038/s41586-020-2698-6 (2020).

70. Ozcan, S. C. *et al.* Stress adaptation pathways and HA–CD44 signaling maintain the survival of pancreatic cancer cells with centrosome amplification. *bioRxiv* DOI: 10.64898/2026.01.24.701523 (2026).

71. Arnandis, T. *et al.* Oxidative stress in cells with extra centrosomes drives non-cell-autonomous invasion. *Dev. Cell* **47**, 409–424, DOI: 10.1016/j.devcel.2018.10.026 (2018).

72. Laughney, A. M., Elizalde, S., Genovese, G. & Bakhoum, S. F. Dynamics of tumor heterogeneity derived from clonal karyotypic evolution. *Cell Reports* **12**, 809–820, DOI: 10.1016/j.celrep.2015.06.065 (2015).

73. Hoshino, A. *et al.* Tumour exosome integrins determine organotropic metastasis. *Nature* **527**, 329–335, DOI: 10.1038/nature15756 (2015).

74. Clements, M. & Johnson, R. Breast Cancer Dormancy in Bone. *Curr. Osteoporos. Reports* **17**, 353–361, DOI: 10.1007/s11914-019-00532-y (2019).

75. Marlow, R. *et al.* A novel model of dormancy for bone metastatic breast cancer cells. *Cancer Res.* **73**, 6886–6899, DOI: 10.1158/0008-5472.CAN-13-0991 (2013).

76. Zhang, J., Cunningham, J. J., Brown, J. S. & Gatenby, R. A. Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nat. Commun.* **8**, 1816, DOI: 10.1038/s41467-017-01968-5 (2017).

77. Gogna, R., Madan, E., Kuppusamy, P. & Pati, U. Re-oxygenation causes hypoxic tumor regression through restoration of p53 wild-type conformation and post-translational modifications. *Cell Death & Dis.* **3**, e286, DOI: 10.1038/cddis.2012.15 (2012).

78. Yang, S. *et al.* Pancreatic cancers require autophagy for tumor growth. *Genes & Dev.* **25**, 717–729, DOI: 10.1101/gad.2016111 (2011).

79. Benci, J. L. *et al.* Opposing functions of interferon coordinate adaptive and innate immune responses to cancer immune checkpoint blockade. *Cell* **178**, 933–948.e14, DOI: 10.1016/j.cell.2019.07.019 (2019).

80. Sahu, D. *et al.* A fluorescent dual-reporter system for real-time monitoring of hif transcriptional activity. *Biochem. Biophys. Res. Commun.* **553**, 59–65, DOI: 10.1016/j.bbrc.2021.03.040 (2021).

81. Kukimoto-Niino, M. *et al.* A fluorescent reporter system for monitoring p21 promoter activity. *Cell Struct. Funct.* **38**, 237–245, DOI: 10.1247/csf.13017 (2013).

82. Harton, M. D. *et al.* p53 pulse dynamics and transcriptional outcomes in single cells. *Mol. Syst. Biol.* **19**, e11270, DOI: 10.15252/msb.202211270 (2023).

83. Enriquez-Navas, P. M. *et al.* Exploiting evolutionary steering to induce collateral drug sensitivity in cancer. *Nat. Commun.* **11**, 1923, DOI: 10.1038/s41467-020-15596-z (2020).

84. Yuan, S., Norgard, R. J. & Stanger, B. Z. Cellular plasticity in cancer. *Cancer Discov.* **9**, 837–851, DOI: 10.1158/2159-8290.CD-19-0015 (2019).

85. Deng, S. *et al.* Tumor cell plasticity in targeted therapy-induced resistance: mechanisms and new strategies. *Signal Transduct. Target. Ther.* **8**, 86, DOI: 10.1038/s41392-023-01383-x (2023).

86. Gatenby, R. A., Silva, A. S., Gillies, R. J. & Frieden, B. R. Adaptive therapy. *Cancer Res.* **69**, 4894–4903, DOI: 10.1158/0008-5472.CAN-08-3658 (2009).

87. Zhang, J., Cunningham, J. J., Brown, J. S. & Gatenby, R. A. Evolution-based mathematical models significantly prolong response to abiraterone in metastatic castrate-resistant prostate cancer and identify strategies to further improve outcomes. *eLife* **11**, e76284, DOI: 10.7554/eLife.76284 (2022).

88. West, J. *et al.* A survey of open questions in adaptive therapy: Bridging mathematics and clinical translation. *eLife* **12**, e84263, DOI: 10.7554/eLife.84263 (2023).

89. Brady-Nicholls, R., Enderling, H. *et al.* Range-bounded adaptive therapy in metastatic prostate cancer. *Cancers* **14**, 5319, DOI: 10.3390/cancers14215319 (2022).

90. Sartakhti, J. S., Manshaei, M. H., Bateni, S. & Archetti, M. Evolutionary dynamics of tumor-stroma interactions in multiple myeloma. *PLoS ONE* **11**, e0168856, DOI: 10.1371/journal.pone.0168856 (2016).

91. Li, Z., Tan, X. & Yu, Y. Optimal adaptive cancer therapy based on evolutionary game theory. *PLoS ONE* **20**, e0320677, DOI: 10.1371/journal.pone.0320677 (2025).

92. Gluzman, M. *et al.* Bringing evolutionary cancer therapy to the clinic: a systems approach. *npj Syst. Biol. Appl.* DOI: 10.1038/s41540-025-00528-8 (2025).

Box 1: A Glossary of Evolutionary Games in Cancer

The game-theoretic concepts used in this article originate from economics and evolutionary biology. Below we define each game type and its cancer-biological interpretation. In all cases, *payoff* refers to the reproductive fitness (net growth rate) of a cell playing a given strategy, and *frequency dependence* means that payoffs change as the proportions of strategies in the population shift.

Nash equilibrium & evolutionarily stable strategy (ESS). A Nash equilibrium is a strategy combination from which no single player can improve its payoff by unilaterally switching. An ESS is the biological counterpart, a strategy that, once dominant, cannot be invaded by any rare mutant. Stable coexistence of proliferative and stress-tolerant strategies, observable as distinct cancer cell states, can represent a mixed-strategy Nash equilibrium maintained by frequency-dependent selection.

Prisoner's Dilemma. Two players each choose to cooperate or defect. Mutual cooperation yields a moderate payoff to both, but each is individually tempted to defect (free-ride) because defection yields the highest payoff against a cooperator. If both defect, both receive a low payoff. In cancer, a pure Prisoner's Dilemma predicts the collapse of cooperation (e.g., growth-factor secretion), because non-producing "cheater" cells always outcompete producers. The persistence of cooperation in tumors therefore implies that additional mechanisms, spatial structure, nonlinear benefits, or repeated interactions, modify the basic dilemma.

Snowdrift (Hawk–Dove) game. Unlike the Prisoner's Dilemma, here a player gains more by cooperating when its opponent defects than by also defecting. This structure guarantees stable coexistence of both strategies. In glioblastoma, glycolytic cells ("hawks" that acidify the microenvironment) and OXPHOS cells ("doves" that conserve glucose) coexist because each gains a frequency-dependent advantage when rare, a hallmark of Snowdrift dynamics.

Public goods game. A multiplayer extension of cooperation dilemmas in which individuals choose how much to invest in a shared resource that benefits all group members. In tumors, secreted factors such as VEGF, TGF-β, and IL-10 function as public goods, wherein producers bear the metabolic cost while all nearby cells, including non-producers, benefit. When the relationship between the number of contributors and the collective benefit is nonlinear (sigmoid), stable coexistence of producers and free-riders is possible.

Tragedy of the commons. The outcome when every individual over-exploits a shared resource, depleting it to the detriment of all. In cancer, therapies that engineer "defector" cells, abolishing growth-factor production while retaining the ability to consume neighbors' secretions, can trigger a tragedy of the commons, collapsing the cooperative infrastructure that sustains collective stress adaptation.

Repeated game. A game played over multiple rounds, allowing players to condition current choices on past outcomes. Tumor evolution unfolds over repeated cycles of division, stress, and therapy. Repeated-game structure enables conditional strategies such as epigenetic memory of prior stress and dynamic equilibria that could not arise in a single round.

Leader game & Deadlock game. In a Leader game, both players prefer to adopt different strategies, but each wants the other to move first; coexistence is an equilibrium. In a Deadlock, mutual defection is the only stable outcome regardless of the opponent's choice. Kaznatcheev *et al.* showed experimentally that NSCLC cells play a Leader game under baseline conditions (enabling sensitive–resistant coexistence) but shift to Deadlock under alectinib treatment, where only resistant cells survive.

Stackelberg (leader–follower) game. A sequential game in which one player (the "leader") commits to a strategy first, and the other (the "follower") responds optimally. In adaptive therapy, the clinician acts as leader, choosing dosing schedules, and the tumor population responds by shifting strategy frequencies. The clinician's challenge is that payoff matrices and strategy frequencies must be inferred from noisy clinical surrogates rather than observed directly.

Metagame. A game about games, a higher-level interaction in which one player's moves alter the rules (payoff matrix) of the game played by others. Therapy is a metagame because chemotherapy, targeted agents, and immunotherapy do not merely select among existing strategies but qualitatively change the type of game being played, for instance, converting a Snowdrift into a Prisoner's Dilemma.

Bet-hedging. A strategy in which individuals stochastically adopt different phenotypes to spread risk across unpredictable environments. In cancer, drug-tolerant persister cells represent a bet-hedging subpopulation that is rare under baseline conditions (switching rate $\sim 10^{-5}$–10^{-4}) and provides insurance against sudden therapeutic stress. Sublethal MOMP and CIN-driven karyotypic diversification extend bet-hedging to the genomic level.

Nested games. Games played simultaneously at multiple hierarchical levels, where the rules at one level constrain strategies at another. In tumors, fast metabolic games (minutes–hours), medium epigenetic games (days–weeks), and slow genetic/CIN games (weeks–months) are nested such that strategies available at each timescale are determined by outcomes at slower timescales.

Asymmetric (coercive) game. A game in which players have different strategy sets or payoff structures. The reverse Warburg effect exemplifies this, as tumor cells coerce CAFs into glycolytic metabolism through paracrine ROS signaling, extracting metabolic subsidies in what resembles parasitism rather than mutual cooperation. The asymmetry creates therapeutic vulnerabilities absent in symmetric interactions.

Box 2: What Current Frameworks Systematically Miss

The dominant conceptual frameworks in cancer biology, pathway activation models, mutation-centric inevitability, and precision oncology's target-and-kill logic, have generated extraordinary mechanistic insight. Yet each contains a systematic blind spot that this framework addresses.

Pathway activation $\neq$ adaptive strategy. Identifying that the DDR, UPR, or HIF pathway is "activated" does not explain why different cells under the same stress adopt different responses, nor why the same pathway activation produces opposite fitness consequences depending on the strategies of neighboring cells. Pathway-centric thinking treats stress responses as deterministic outputs of molecular circuits; the game-theoretic framework treats them as frequency-dependent strategic choices whose fitness consequences depend on population composition.

Mutation-centric inevitability obscures reversibility. The standard narrative, mutations accumulate, resistance is inevitable, systematically underestimates the role of reversible, epigenetically mediated strategy switching. Drug-tolerant persister states, metabolic reprogramming, and immune evasion programs arise and dissolve on timescales far shorter than mutational acquisition. Treating resistance as genetically determined leads to therapeutic nihilism; recognizing it as a game-theoretic equilibrium reveals that it can be destabilized.

Treating heterogeneity as noise blocks prediction. Intratumor heterogeneity is routinely described as a barrier to treatment rather than as a structured, predictable outcome of evolution under constraint. The game-theoretic framework reframes heterogeneity as the expected equilibrium of frequency-dependent selection, maintained because no single strategy can dominate across all conditions, and thereby converts it from an obstacle into a source of testable predictions about tumor behavior and therapeutic vulnerability.

Box 3: Testable Predictions from the Game-Theoretic Framework

1. Heterogeneity scales with stress variability. If heterogeneity reflects evolutionarily stable strategy mixtures, then cell state entropy (measured by scRNA-seq as a proxy for strategy diversity) should correlate with microenvironmental variability (perfusion fluctuation, immune infiltration heterogeneity). Intermediate stress variability should maximize phenotypic diversity; very stable or very chaotic environments should favor monomorphic populations.

2. Plasticity rates match fluctuation timescales. If plasticity is evolved bet-hedging, switching rates between cell states (reflecting underlying strategy transitions driven by changing game states) should scale with the frequency of environmental fluctuations. *Test*: Compare phenotypic switching rates (live-cell imaging or lineage tracing) under controlled cycling vs. constant stress, and compare observed rates to predictions from bet-hedging models parameterized with empirical stress dynamics.

3. Collateral sensitivities follow from trade-off structure. Adaptation to one stress should create predictable vulnerabilities to mechanistically opposed stresses. *Test*: Expose cells to a stress panel, then measure drug sensitivity across a complementary panel. The resulting stress-sensitivity matrix should exhibit anti-correlated payoff vectors along axes predicted by known trade-offs (glycolysis vs. OXPHOS, proliferation vs. survival, biosynthesis vs. autophagy).

4. Spatial structure stabilizes cooperation. Cooperative phenotypes (VEGF secretion, TGF-β production) should persist at higher frequencies in spatially structured regions than in well-mixed regions. *Test*: Compare cooperator frequency between compartmentalized tumor-stroma interfaces and less structured regions using spatial transcriptomics; manipulate spatial structure in organoid or microfluidic models and measure cooperator stability.

5. Adaptive therapy efficacy depends on interaction structure. Adaptive dosing benefits should scale with the strength of competitive suppression of resistant by sensitive cells. *Test*: Measure competitive fitness in co-culture at varying ratios, parameterize game-theoretic models, and predict optimal dosing schedules. Strong competition should yield large adaptive therapy benefit; weak or mutualistic interactions should yield minimal benefit.

6. Pre-existing strategy diversity accelerates adaptation to novel stress. Tumors with high baseline cell state diversity (as a proxy for strategy diversity) should reach resistance (defined as growth rate recovery $\geq$50% of pre-treatment levels) faster than low-diversity tumors when exposed to drugs outside their selective history. *Test*: Measure baseline cell state entropy across a cell line panel via scRNA-seq, expose each to a novel drug, and correlate time-to-resistance with pre-treatment entropy. Diversity specifically in stress-response gene modules (not housekeeping heterogeneity) should predict adaptive speed, and spatial profiling should confirm that cell state diversity maps onto distinct game states and strategies rather than reflecting stochastic noise within a single game state.