Conservation Biology

Conservation Practice and Policy

The costs of avoiding environmental impacts from shale-gas surface infrastructure

Austin W. Milt,* ¶ Tamara D. Gagnolet,† and Paul R. Armsworth*

*Department of Ecology & Evolutionary Biology, University of Tennessee, Knoxville, 569 Dabney Hall, 1416 Circle Drive, Knoxville, TN 37996, U.S.A.

†The Nature Conservancy, Pennsylvania Chapter, 2101 North Front Street, Building #1, Suite 200, Harrisburg, PA 17110, U.S.A.

Abstract: *Growing energy demand has increased the need to manage conflicts between energy production and the environment. As an example, shale-gas extraction requires substantial surface infrastructure, which fragments habitats, erodes soils, degrades freshwater systems, and displaces rare species. Strategic planning of shale-gas infrastructure can reduce trade-offs between economic and environmental objectives, but the specific nature of these trade-offs is not known. We estimated the cost of avoiding impacts from land-use change on forests, wetlands, rare species, and streams from shale-energy development within leaseholds. We created software for optimally siting shale-gas surface infrastructure to minimize its environmental impacts at reasonable construction cost. We visually assessed sites before infrastructure optimization to test whether such inspection could be used to predict whether impacts could be avoided at the site. On average, up to 38% of aggregate environmental impacts of infrastructure could be avoided for 20% greater development costs by spatially optimizing infrastructure. However, we found trade-offs between environmental impacts and costs among sites. In visual inspections, we often distinguished between sites that could be developed to avoid impacts at relatively low cost (29%) and those that could not (20%). Reductions in a metric of aggregate environmental impact could be largely attributed to potential displacement of rare species, sedimentation, and forest fragmentation. Planners and regulators can estimate and use heterogeneous trade-offs among development sites to create industry-wide improvements in environmental performance and do so at reasonable costs by, for example, leveraging low-cost avoidance of impacts at some sites to offset others. This could require substantial effort, but the results and software we provide can facilitate the process.*

Keywords: access roads, conservation planning, environmental externalities, gathering pipelines, multiobjective planning, shale-energy policy, spatial optimization, well pads

Los Costos de Evitar los Impactos Ambientales de la Infraestructura Superficial del Gas de Esquisto

Resumen: *La creciente demanda de energía ha incrementado la necesidad de manejar los conflictos entre la producción de energía y el ambiente. Como ejemplo, la extracción de gas esquisto requiere de una infraestructura superficial sustancial, la cual fragmenta los hábitats, erosiona el suelo, degrada los sistemas de agua dulce y desplaza a las especies raras. La planeación estratégica de la infraestructura de gas esquisto puede reducir las compensaciones entre los objetivos económicos y ambientales, pero la naturaleza específica de estas compensaciones no se conoce. Estimamos el costo de evitar los impactos del cambio de uso de suelo causado por el desarrollo de gas esquisto dentro de los arriendos sobre los bosques, humedales, especies raras y arroyos Creamos un software para sitiar óptimamente la infraestructura superficial de gas esquisto y minimizar su impacto ambiental a un costo de construcción razonable. Valoramos visualmente los sitios antes de la optimización de la infraestructura para probar si dicha inspección podría usarse para predecir si los impactos podrían evitarse en el sitio. En promedio, hasta el 38 % de los impactos ambientales agregados de la infraestructura podría evitarse por 20 % de costos de desarrollo mayores al optimizar espacialmente la infraestructura. Sin embargo, encontramos compensaciones entre los impactos ambientales y los costos entre los sitios. En las inspecciones visuales muchas veces distinguimos entre los sitios que podrían desarrollarse para evitar los impactos a un costo relativamente bajo (29 %) y aquellos que no podrían (20 %). Las reducciones en*

¶*email milt@wisc.edu*
Paper submitted March 22, 2016; revised manuscript accepted May 24, 2016.

Conservation Biology, Volume 00, No. 0, 1–8

DOI: 10.1111/cobi.12766

una medida de impacto ambiental agregado podrían atribuirse en su mayoría al desplazamiento potencial de las especies raras, la sedimentación y la fragmentación del bosque. Los planificadores y los reguladores pueden estimar y usar compensaciones heterogéneas entre los sitios de desarrollo para crear mejoras en el desempeño ambiental a lo largo de la industria y hacerlo a costos razonables al, por ejemplo, evitar los impactos en algunos sitios para compensar otros. Esto podría requerir un esfuerzo sustancial, pero los resultados y el software que proporcionamos pueden facilitar el proceso.

Palabras Clave: agrupación de tuberías, carreteras de acceso, conjunto de pozos, externalidades ambientales, optimización espacial, planeación de la conservación, planeación multiobjetiva, políticas de gas de esquisto

Introduction

As countries increase domestic energy production, alternative energy sources and associated environmental effects are increasingly scrutinized. Production of shale gas in the United States has grown markedly since 2008, and in 2012 it was the largest source of natural gas (U.S. Energy Information Administration 2014). Other countries, for example, the United Kingdom (Hays et al. 2015) and China, are debating their own gas development. Research increasingly highlights environmental, human health, social, and economic effects of shale-gas production (Brittingham et al. 2014; Evans & Kiesecker 2014; Hays et al. 2015). Surface infrastructure can negatively affect terrestrial and freshwater biodiversity through habitat loss and fragmentation (Gillen & Kiviat 2012; Kiviat 2013; Jones et al. 2014), pollution, and sedimentation (Kassotis et al. 2013; Olmstead et al. 2013).

As a topic, shale energy has been added to ongoing (Jenner & Lamadrid 2013), decades-long discussions about the impacts of energy production (e.g., Dincer 1999; Demirbas 2009; Arvesen & Hertwich 2012). We focused on the surface terrestrial and freshwater impacts of well pads, access roads, and gathering pipelines required for directional drilling and hydraulic fracturing at the leasehold level. A leasehold is formed by aggregating up to tens of adjacent subsurface development rights into a single planning unit that is many thousands of hectares. The footprint of shale-gas construction sites (Fig. 1) varies, but one well pad with its associated infrastructure has a footprint similar to that of a rural home. Well pads are gravel stages (approximately 0.01 km^2) for well heads - the point at which the vertical portion of a gas well meets the surface - and additional equipment needed to process gas before transport. At such sites, well pads are separated by hundreds of meters. Similarly, in half the counties in Pennsylvania (U.S.A.)—our case study area—homes are on average >180 m apart (U.S. Census Bureau 2016). Access roads, which connect well pads to existing roads, are short (approximately 0.1 km here), narrow (12 m here), and made of gravel. Half of rural roads in Pennsylvania are <0.25 km long (2015 TIGER/Line Shapefiles prepared by U.S. Census Bureau, 2015). Gathering pipelines, which transport gas offsite, are similar in corridor size and straightness to buried electrical transmission lines.

Shale-gas developers often plan infrastructure configurations by hand (Triana Energy LLC, personal communication). Optimization can be used to explore feasible configurations that reduce potential environmental impacts relatively inexpensively. We developed software that examines the spatial optimization problem of siting well pads, access roads, and gathering pipelines to minimize environmental impacts, including and beyond regulatory requirements. Solutions in planning projects with multiple spatial objectives often have trade-offs (e.g., Raudsepp-Hearne et al. 2010; Qiu & Turner 2013; Ruijs et al. 2013), which we sought to address. Infrastructure planners try to minimize costs, but conservation-oriented planning may increase costs. Therefore, a vital first step to making informed progress is to quantify the costs of avoiding environmental impacts beyond regulatory requirements. We concentrated on major costs that vary with the spatial configuration of infrastructure: moving earth, clearing land, constructing stream-crossing infrastructure, and materials and labor (Triana Energy LLC, personal communication).

To the best of our knowledge, we are the first to explicitly optimize the planning of well pads, access roads, and gathering pipelines to minimize aggregate impacts of shale-gas development at reasonable construction costs. Others have quantified the environmental impacts of shale-gas surface infrastructure (Davis & Robinson 2012; Olmstead et al. 2013; Evans & Kiesecker 2014; Racicot et al. 2014; Slonecker & Milheim 2015), but no one has examined trade-offs between costs of surface infrastructure development and its environmental impacts. Costs of avoiding impacts have been given limited consideration at best (Keller et al. 2015).

Existing software would not have adequately served our purpose. Most well-known spatial conservation planning software designates protection to decision units of a homogenous type to increase benefits for a target (Moilanen et al. 2009). A somewhat recent advance incorporates corridors in planning software by adding rules for site selection (e.g., Pouzols & Moilanen 2014). Our software integrates siting decisions for multiple infrastructure types of differing geometry, each of which

Figure 1. Typical shale-gas surface infrastructure in Pennsylvania (U.S.A.), the region of our case study. Well pads (rectangular clearings) are stages for equipment and well bores, access roads (linear clearings in bottom right) give vehicle access to a site, and gathering pipelines (other linear clearings) transport gas from a site. Photo by M. Godfrey.

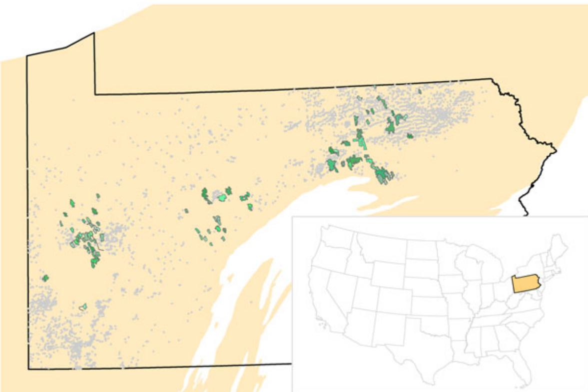

Figure 2. Shale-gas development sites (polygons) and grouped well permits (gray dots) in the Marcellus shale formation (shading), where shale-gas development has been concentrated in Pennsylvania (U.S.A.).

differentially affects planning objectives. The oil and gas industry also lacks software that simultaneously plans multiple types of infrastructure with environmental impact avoidance as a primary objective.

We aimed primarily to quantify site-level costs of avoiding environmental impacts from the land-use change associated with shale-gas development. We used shale-gas surface-infrastructure planning in Pennsylvania as a case study. In particular, we developed a novel, advanced spatial optimization algorithm to plan well-pad locations and access road and pipeline routes at 84 sites in 5 counties (Fig. 2) and quantified the trade-off between avoiding environmental impacts and increasing construction costs. When the results of our analysis revealed overall promising but heterogeneous potential to avoid impacts at reasonable costs, we further explored causes for site-level differences in cost-impact trade-offs.

Methods

Bungee Infrastructure Planning Software

Bungee is the infrastructure planning software we devised (details in Supporting Information). Bungee stands for balancing unconventional natural gas extraction and the environment. It can be used to plan surface infrastructure in leaseholds (Supporting Information) while minimizing environmental impacts from infrastructure beyond regulatory requirements and with a limited construction budget. Bungee is programmed in Python 2.7 (Python Software Foundation, https://www.python.org) and Cython 0.19 (cython.org). It has 2 major parts: the Python module, which is used like any other Python module, and Bungee GIS, a Python toolbox for ArcGIS. A streamlined interface to Bungee is provided by The Nature Conservancy in a Python toolbox called Landscape Environmental Energy Planning (LEEP). Both require ArcGIS 10.1+ with the Spatial Analyst Extension. Bungee is like ArcGIS's arcpy module: major workflow tools and other minor tools are all accessible within Python. Bungee analyses take hours to days depending on leasehold size. For instance, using a typical desktop computer, Bungee placed infrastructure in a 12-km^2 area in 45 min, but took several days in a 72-km^2 area. Future versions could use lower level languages and proprietary solvers to improve runtime. Both Bungee and LEEP can be obtained from the authors (http://www.austinmilt.com/bungee).

Bungee has 3 major conceptual steps. First, determine where and how much infrastructure may be placed. Second, prepare data to assess the environmental impact of proposed infrastructure. Third, plan infrastructure layouts–configurations of well pads, access roads, and gathering pipelines within a site.

First, Bungee prepares planning constraints based on default settings and user inputs, including regulatory setbacks and construction practices. Regulatory setbacks, for example, minimum distances from buildings, define off-limit areas for development. Construction practices also constrain infrastructure. For instance, well-pad locations are constrained by their production units (the geometric area of gas extracted at a well pad). In our study area, a production unit was 6 parallel wells, 3 on either side of the well pad. Bungee iteratively adds nonoverlapping production units to the leasehold while ensuring none extend beyond the leasehold boundary by more than a tolerance (20% here) of production-unit area. In each iteration, simulated annealing shuffles production units to make room for additional units (Supporting Information). This optimization maximally drains the leasehold of gas while minimizing the number of well pads. Another example of a construction practice is the construction budget, which in our case study was double the cost of a cost-minimizing layout.

Second, Bungee prepares inputs for the optimization objective function. Bungee minimizes an aggregate environmental impact score subject to a budget constraint. In practice, the objective and budget constraint are combined into a synthetic objective function at runtime to analyze trade-offs between the 2. A layout's impact score is a weighted sum of normalized impact metrics specified by the user (see "Case Study Details" below & Supporting Information). A layout's impact score is a weighted sum of normalized impact metrics specified by the user (see "Case Study Details" below & Supporting Information). Weights can be adjusted to tailor Bungee analyses to local priorities. Weights could come from valuation exercises or other studies (e.g., Banzhaf et al. 2016). Each metric's normalization constant is its value for its minimizing layout (Supporting Information for case-study specifics). The average construction cost across layouts from this step is used as a normalization constant for costs.

Bungee uses a modified form of Dijkstra's least cost path algorithm to plan linear infrastructure (Dijkstra 1959). We added a secondary construction-cost surface to invalidate otherwise environmentally optimal routes. The least cost algorithm in Bungee uses a set of strictly positive, spatially additive (i.e., effects can be directly summed over space) impact surfaces, which are derived for each metric and infrastructure type by calculating the impact score for each pixel in the leasehold as if that were the only pixel developed.

Third, Bungee plans infrastructure layouts with a genetic algorithm optimization (Yu & Gen 2010) (Supporting Information). Dijkstra's algorithm finds the single optimum route from a point or points to one or more destinations. Consequently, linear infrastructure routes depend only on the locations of well pads and existing infrastructure. This allows one to represent layouts by the coordinates of well pads and the order in which linear infrastructure is planned. Bungee allows linear infrastructure to connect to planned infrastructure such that planned infrastructure affects the routes of subsequent infrastructure. Bungee starts with a randomly drawn population of (here $n = 20$) layouts. Low-impact pairs of layouts are preferentially chosen, and the pair's attributes are randomly combined to form a new candidate layout. Each new layout is altered by offsetting pad locations and shuffling the order of linear infrastructure. Linear infrastructure is then planned using our modified Dijkstra's algorithm (Supporting Information). Lower impact new layouts are kept for the subsequent round of the algorithm, and higher-impact new layouts are kept with some probability. In this way, Bungee iteratively improves the population of layouts until termination criteria are met. Bungee repeats this genetic algorithm optimization with an incrementally increasing construction budget, returning an optimal layout at each budget, and finally returning a set of layouts that traces the trade-off curve of impact and cost.

Case Study

We used Pennsylvania as a representative area for shale-energy development in the eastern United States. Over 9620 horizontal wells were drilled in Pennsylvania from 2008 to 2014 (Supporting Information), and Pennsylvania has high levels of biodiversity (Stein et al. 2000).

Shale-gas lease boundaries are not publicly available, so we demarcated sites based on locations of unconventional, horizontal, active, or regulatory inactive wells drilled from 2008 to 2013 (Supporting Information). Well-pad locations were approximated by grouping nearby well points. Each well-pad point was overlaid with a 6-well production unit that was 3353 m tall by 914 m wide ($3000 \times 11{,}000$ feet) and rotated 27° counterclockwise (Triana Energy LLC, personal communication). This rotation is representative of local gas operations. We overlaid production units on land parcels. A set of contiguous land parcels shared by production units of a single operator became one site. When combined with sufficient access to existing pipeline and road infrastructure, this process produced 84 developed sites ranging in size from 465 to 7297 ha (Fig. 2).

Five metrics formed the impact score: forest area lost (forest loss) through development of forest pixels; total edge-to-area ratio of forest after construction (forest frag.) as a measure of forest fragmentation; wetland encroachment (wetlands) as the percentage of a 61–91 m (200–300 feet) area around wetlands occupied by infrastructure; potential sedimentation in water bodies (sediment) (Supporting Information); and expected impact on rare species (rare spp.) as the expected number of known rare species occurrences overlapped by infrastructure based on habitat associations across the state. Any publicly available software approach that uses publicly available rare species inventories to assess impacts on rare species will have limited efficacy because rare species inventories are relatively sparse and those data are often protected to avoid their use for illicit purposes. Nonetheless, the rare spp. metric should in general help direct infrastructure away from areas associated with known rare species occurrences. In our case study, we weighted each category (forests, wetlands, streams, and rare species) of features equally, such that forest loss and forest frag. were each half the weight of others.

As with many popular conservation planning tools that attempt to embrace real-world complexity, Bungee's optimization may produce local (rather than global) optima. Consequently, for each site, we ran optimized infrastructure 5 times at 40 budget increments each time (200 optimizations) and retained those layouts that did not simultaneously have more of an impact and were

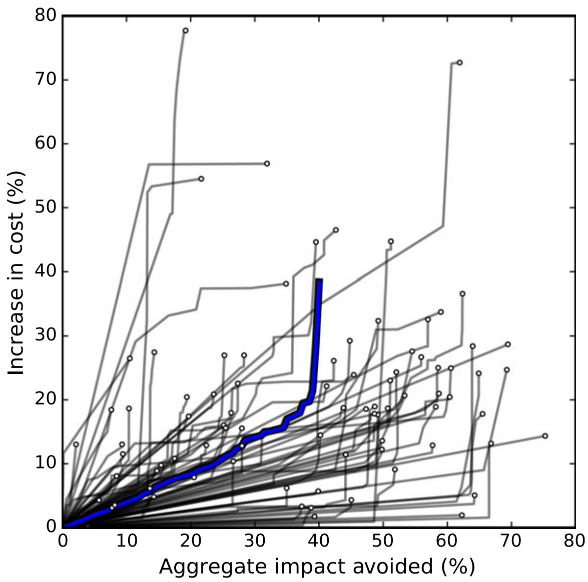

Figure 3. Cost of avoiding impacts of surface infrastructure at the development-area scale as the commitment to reducing impact increases. Thin lines are trade-off curves for individual sites, illustrating the general shape of trade-off curves, the maximum avoidance of aggregate environmental impacts, and the resulting increase in construction costs. Lines were interpolated between individual layouts for clarity (in reality, they are discrete points). Open circles are endpoints, shown to facilitate interpretation. The thick line is the median of the trade-off curves, where truncated curves are given a very high cost at higher levels of impact avoidance.

more costly than any other layout (i.e., they were more Pareto optimal than those not retained).

To inform future policy and practice, we sought to predict site-level trade-offs without planning. We overlaid layouts on a raster approximation of the impact score (Supporting Information) and then visually compared the locations of high-impact layouts with lower impact layouts relative to the underlying raster. Subsequently, we built regression models in an attempt to predict site-level capacity to avoid impacts based on site attributes.

Results

Cost of Avoiding Impacts

On average across our study sites, avoiding 38% of the aggregate environmental impact (impact) of development increased development costs by 20% ($0.8–3.8 million/well pad) (Fig. 3). Further avoidance beyond 38% quickly became cost prohibitive (Fig. 3). Trade-off curves were linear interpolated between (impact and cost) points for each layout relative to the least cost layout for a site. The median line can be thought of as the outcome if using a median-based uniform action across sites.

Large, low-cost reductions of potential aggregate impacts (Fig. 4) were possible for some sites, but the results were highly variable (Fig. 3). One site avoided up to 75% of impacts for $4.1 million (+15%, $1 million/well pad). Another site avoided 2% of impacts at a cost of $7.7 million (+14%, $1.1 million/pad). Most sites (75 of 84 or 89%) had trade-off curves closer to the horizontal than vertical axis, indicating that it was relatively cheap to avoid impacts in those sites (Fig. 3). We expected nonconvex curves. Sites exhibited multiple classes of feasible layouts with different spatial patterns, costs, and impacts. For example, at 1 site, 1 class had 3 disjoint roads, whereas another had 4. Each class had a unique, convex trade-off curve, and when Pareto-optimal points along these curves were sampled by Bungee, the resulting trade-off curve was not always convex. This will be true generally.

Infrastructure Layouts

Bungee automatically produced a variable number of layouts at each site ($n = 2$–20), even though within a site all layouts were constrained to have the same number of well pads and wells per pad ($n = 6$). Many sites had a similar spatial structure of infrastructure layouts (Fig. 4). At low costs, access roads and pipelines tended to be straight (Figs. 4a & 4c) and short. Larger budgets resulted in layouts that tracked lower impact areas (Figs. 4b & 4d). These differences between layouts were often subtle but in some cases greatly affected impacts and costs. For instance, in Fig. 4, the differences in layouts between Figs. 4a and 4b increased costs by $1.5 million and avoided impacts to all but wetlands.

Avoiding Impact at a Low Cost

In most cases, visual inspection differentiated between sites (25 of 84 or 29%) that if used would avoid much impact for little cost (i.e., those with trade-off curves in the lower right of Fig. 3) and sites (17 of 84 or 20%) that if used would not avoid impact. Two conditions distinguished the aforementioned groups: whether the least cost layout had at least some infrastructure in high-impact areas and whether there were lower impact areas for infrastructure to be placed. The 25 sites where the least cost layout had at least some infrastructure in high-impact areas tended to have planned roads or pipelines in high-impact areas (21 of 25 sites). Much less often (4 of 25 sites) planned well pads, but not roads or pipelines, were in high-impact areas. Those 17 sites in the second group, associated with either no large or very costly impact avoidance, tended to be constrained by impacts. Nine of 17 sites lacked low-impact alternatives for pipelines and pads. Often, sparse existing pipelines in isolation

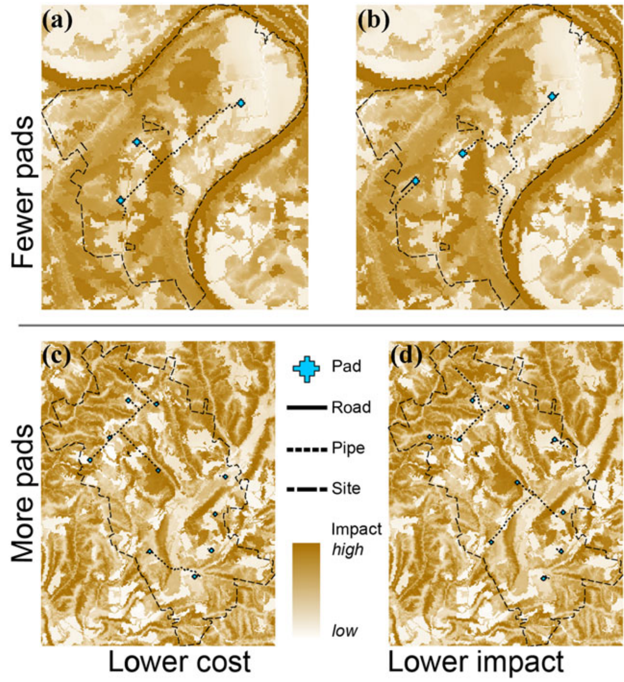

Metric	Layout a	Layout b
Impact score	2.08	0.78
Construction cost ($M)	3.9	5.4
Cost per pad ($M)	1.3	1.8
Rare spp. (spp. × 10^{-3})	0.57	0.00
Wetlands (%)	0.01	0.11
Wediment (t·y^{-1})	13.4	0.1
Forest loss (ha)	2.0	1.1
Forest frag. (k·km^{-2})	18.7	18.1

Metric	Layout c	Layout d
Impact score	2.16	0.90
Construction cost ($M)	12.1	14.9
Cost per pad ($M)	1.2	1.5
Rare spp. (spp. × 10^{-3})	1.09	0.14
Wetlands (%)	0.05	0.07
Sediment (t·y^{-1})	124.9	0.0
Forest loss (ha)	6.6	3.9
Forest frag. (k·km^{-2})	13.6	14.1

Figure 4. Example infrastructure layouts and their estimated impacts produced with our infrastructure planning software (see methods & Supporting Information). The choice of layouts was made to illustrate how reducing potential impacts and increasing construction costs lead to differences in spatial configuration. Individual impact metrics (Supporting Information) are rare spp. (expected number of known rare species locations encountered), wetlands (percentage of area around wetlands occupied by infrastructure), sediment (sediment load in moving water bodies caused by soil disturbance), forest loss (area of forest cleared), and forest frag (forest edge-to-area ratio after development). Impact score is an aggregate of the individual impact metrics and is not comparable among sites. We used a monotonic, nonlinear color scaling on the impact surface to enhance heterogeneity of impact values.

or along with feasible well-pad locations forced gathering pipelines through high-impact areas. In other sites (6 of 17), reductions in some impacts led to increases in others, such that the aggregate environmental impacts did not change much. In 3 of 17 sites, the least cost layout was already in a low-impact area, preventing further impact avoidance. Our regression models revealed little predictive capacity, such that easily observed site attributes such as site area, number of well pads, slope, correlation between impact and cost surfaces, and existing road and pipeline density showed little association with our responses (Supporting Information).

Metrics Driving Impact Avoidance

Individual impact metrics differentially contributed to aggregate impact avoidance relative to the least cost layout (Supporting Information). Reductions in the rare spp. metric contributed to a mean of 40% (SE 5) of aggregate impact avoidance across sites, followed by sediment yield (25% [SE 5]), forest loss (19% [3]), forest frag. (−0.1% [0.2]), and wetlands (−1.2% [0.8]). We note the relative contribution of each metric depended partially on its weight in the optimization. Variable contributions of impact metrics to aggregate impact avoidance were due to combinations of some metrics having high spatial variance over small scales, high correlation with cost such that low-cost options were also low impact, or trade-offs with other metrics.

Discussion

As energy development continues worldwide, opportunities exist to reduce potential impacts through conservation-oriented planning. We quantified the monetary cost of such planning at the site level for shale-gas development in Pennsylvania. On average, sites can avoid about 38% of aggregate environmental impacts at a 20% increased construction cost (Fig. 3), but further impact avoidance is very expensive. Impacts could be avoided at many sites for low cost (Fig. 3 curves extending to lower right). Visual inspection of a least cost layout overlaid on an impact surface was better for predicting site-level impact and cost trade-offs than regression. For instance, in 21 sites, we qualitatively predicted that because some linear infrastructure in the least cost layout was routed through high-impact areas and some low-impact areas were present, much impact could be avoided for relatively low cost. Some impact metrics contributed more than others to avoiding impact. For example, 40% of aggregate impact avoidance was attributable to potential impact on rare species.

Cost-impact trade-offs were heterogeneous across sites, which should inform policies and practices affecting site-level planning. There is scope to avoid more than one-third of impact for a policy that holds developers to the standard of a median-impact site. This is thematically similar to existing environmental shale-gas regulations, which impose uniform constraints on development. Although this policy would increase median costs by up to $400,000/well pad (+20%) (Fig. 3), it may be a small cost relative to total construction costs (Hefley & Seydor 2015). A more flexible regulatory alternative, such as a market-based mechanism, could exploit heterogeneity across sites, incentivizing compensation of companies that choose to impact less by those that impact more. Many other efficient policy options exist (Ferraro 2008).

The choice of environmental impact metrics used in planning affects which actions can be implemented efficiently. Actions targeting one impact may effectively reduce aggregate impacts. We found when assuming equal impact weightings that approximately 65% of reductions in the impact score across sites over the least cost baseline were attributable to just 2 impact metrics related to rare species' habitats and freshwater sedimentation (Supporting Information). Very rarely did impacts trade-off with one another. These results suggest that Pennsylvania could further protect habitat of rare species and areas with high degrees of slope to avoid aggregate impacts across sites. This would be less flexible than targeting aggregate impacts and thus could preclude development in places where those single impacts could not be avoided. Further, a general regulatory tool for aggregate impacts could work in multiple regions where regulating single, location-specific impacts might not. Regardless of the chosen metrics, determining the feasibility of avoiding impacts requires cost-impact trade-off assessment. Results of our qualitative analysis suggest that it may be possible for human or machine-learning classifiers to assess potential trade-offs in sites without sophisticated infrastructure planning and that this may be more informative than typical regression. We chose visual inspection and typical regression because of their ease of implementation and interpretation and their reflection of current decision-making practices. Machine-learning may be able to assimilate more complex relationships and situational variables than either of the others and produce quantitative predictions.

Our study is a novel contribution to the shale-energy environmental impact and policy literature, and Bungee is a novel planning software in multiple regards. We suspect the following can be generalized from our case study: it is possible to avoid much impact inexpensively in some places; sites within a region will have heterogeneous impact-cost trade-offs; layouts forming trade-off curves will be in discrete and spaced groups; and some impacts will contribute more than others to avoid aggregate impacts. Specific results will change with planning characteristics, including the spatial distribution and regulatory control of gas rights, relevant environmental impacts, and developer practices. Bungee can accommodate many variations in planning practices. Consequently, our study represents a step toward a new class of conservation planning problems and solutions: the planning of development to minimize impacts instead of (or in addition to) preventing development in priority areas. Our results can uniquely inform policies and actions at the scale at which development decisions are made. Findings at this scale should also inform decisions at larger scales, for example by influencing regionally applied policies as well as industry choices for when, where, and how to develop.

Acknowledgments

We gratefully acknowledge the feedback of C. Dumoulin, R. Fovargue, G. Iacona, E. Larson, D. Le Bouille, N. Sutton, and S. Ward. We are especially grateful to Triana Energy LLC for providing information and feedback with regard to the costs and process of shale-gas surface infrastructure planning. We appreciate feedback and application from current and future users of the software we developed as part of this research. This work was conducted by AWM while a graduate research assistant at the National Institute for Mathematical and Biological Synthesis, an institute sponsored by the National Science Foundation through NSF award DBI-1300426. Additional support was provided by The University of Tennessee, Knoxville, and The Nature Conservancy.

Supporting Information

Additional methods (including Bungee methods, descriptions of impact metrics, case study data and settings for Bungee analyses, and regressions) and results (Appendix S1) are available online. The authors are solely responsible for the content and functionality of these materials. Queries (other than absence of the material) should be directed to the corresponding author.

Literature Cited

Arvesen A, Hertwich EG. 2012. Assessing the life cycle environmental impacts of wind power: a review of present knowledge and research needs. Renewable and Sustainable Energy Reviews **16:**5994–6006.

Banzhaf HS, Burtraw D, Chung Criscimagna S, Cosby BJ, Evans DA, Krupnick AJ, Siikamäki JV. 2016. Policy analysis: valuation of ecosystem services in the southern Appalachian Mountains. Environmental Science & Technology **50:**2830–2836.

Brittingham MC, Maloney KO, Farag AM, Harper DD, Bowen ZH. 2014. Ecological risks of shale oil and gas development to wildlife, aquatic resources and their habitats. Environmental Science & Technology **48:**11034–11047.

Davis JB, Robinson GR. 2012. A geographic model to assess and limit cumulative ecological degradation from Marcellus Shale Exploitation in New York, USA. Ecology and Society **17:**25. DOI: 10.5751/ES-04822-170225.

Demirbas A. 2009. Political, economic and environmental impacts of biofuels: a review. Applied Energy **86:**S108–S117.

Dijkstra EW. 1959. A note on two problems in connexion with graphs. Numerische Mathematik **1:**269–271.

Dincer I. 1999. Environmental impacts of energy. Energy Policy **27:**845–854.

Evans JS, Kiesecker JM. 2014. Shale gas, wind and water: assessing the potential cumulative impacts of energy development on ecosystem services within the Marcellus Play. PLOS ONE **9** (e89210) DOI: dx.doi.org/10.1371/journal.pone.0089210.

Ferraro PJ. 2008. Asymmetric information and contract design for payments for environmental services. Ecological Economics **65:**810–821.

Gillen JL, Kiviat E. 2012. Hydraulic fracturing threats to species with restricted geographic ranges in the eastern United States. Environmental Practice **14:**320–331.

Hays J, Finkel ML, Depledge M, Law A, Shonkoff SBC. 2015. Considerations for the development of shale gas in the United Kingdom. The Science of the Total Environment **512–513C:**36–42.

Hefley W, Seydor S. 2015. Direct economic impact of the value chain of a Marcellus shale gas well. Pages 15–45 in Hefley WE, Wang Y, editors. Economics of unconventional shale gas development SE—2. Springer International Publishing.

Jenner S, Lamadrid AJ. 2013. Shale gas vs. coal: policy implications from environmental impact comparisons of shale gas, conventional gas, and coal on air, water, and land in the United States. Energy Policy **53:**442–453.

Jones IL, Bull JW, Milner-Gulland EJ, Esipov AV, Suttle KB. 2014. Quantifying habitat impacts of natural gas infrastructure to facilitate biodiversity offsetting. Ecology and Evolution **4:**79–90.

Kassotis CD, Tillitt DE, Davis JW, Hormann AM, Nagel SC. 2013. Estrogen and androgen receptor activities of hydraulic fracturing chemicals and surface and ground water in a drilling-dense region. Endocrinology **155:**897–907.

Keller AA, Fournier E, Fox J. 2015. Minimizing impacts of land use change on ecosystem services using multi-criteria heuristic analysis. Journal of Environmental Management **156:**23–30.

Kiviat E. 2013. Risks to biodiversity from hydraulic fracturing for natural gas in the Marcellus and Utica shales. Annals of the New York Academy of Sciences **1286:**1–14.

Moilanen A, Wilson KA, Possingham HP. 2009. Spatial conservation prioritization: quantitative methods and computational tools. Oxford University Press, New York.

Olmstead SM, Muehlenbachs LA, Shih J-S, Chu Z, Krupnick AJ. 2013. Shale gas development impacts on surface water quality in Pennsylvania. Proceedings of the National Academy of Sciences of the United States of America **110:**4962–4967.

Pouzols FM, Moilanen A. 2014. A method for building corridors in spatial conservation prioritization. Landscape Ecology **29:**789–801.

Qiu J, Turner MG. 2013. Spatial interactions among ecosystem services in an urbanizing agricultural watershed. Proceedings of the National Academy of Sciences of the United States of America **110:**12149–12154.

Racicot A, Babin-Roussel V, Dauphinais J-F, Joly J-S, Noël P, Lavoie C. 2014. A framework to predict the impacts of shale gas infrastructures on the forest fragmentation of an agroforest region. Environmental Management **53:**1023–1033.

Raudsepp-Hearne C, Peterson GD, Bennett EM. 2010. Ecosystem service bundles for analyzing tradeoffs in diverse landscapes. Proceedings of the National Academy of Sciences of the United States of America **107:**5242–5247.

Ruijs A, Wossink A, Kortelainen M, Alkemade R, Schulp CJE. 2013. Trade-off analysis of ecosystem services in Eastern Europe. Ecosystem Services **4:**82–94.

Slonecker ET, Milheim LE. 2015. Landscape disturbance from unconventional and conventional oil and gas development in the Marcellus Shale Region of Pennsylvania, USA. Environments **2:**200–220.

Stein BA, Kutner LS, Adams JS, editors. 2000. Precious heritage: the status of biodiversity in the United States. Oxford University Press, New York.

U.S. Census Bureau. 2016. 2010 U.S. Census housing densities by county. U.S. Census Bureau, Washington, D.C. Available from http://factfinder.census.gov/ (accessed February 2016).

U.S. Energy Information Administration. 2014. Annual energy outlook 2014. U.S. Energy Information Administration, Washington, D.C. Available from http://www.eia.gov/forecasts/aeo/mt_naturalgas.cfm (accessed August 31, 2016).

Yu X, Gen M. 2010. Introduction to evolutionary algorithms. Springer Science & Business Media, New York.